

PAPAL: A Provable PArticle-based Primal-Dual ALgorithm for Mixed Nash Equilibrium

Shihong Ding*

Peking University

DINGSHIHONG@STU.PKU.EDU.CN

Hanze Dong*†

Salesforce AI Research

HENDRYDONG@GMAIL.COM

Cong Fang ✉

Peking University

FANGCONG@PKU.EDU.CN

Zhouchen Lin

Peking University

ZLIN@PKU.EDU.CN

Tong Zhang

University of Illinois Urbana-Champaign

TONGZHANG@TONGZHANG-ML.ORG

Editor: Francesco Orabona

Abstract

We consider the non-convex non-concave objective function in two-player zero-sum continuous games. The existence of pure Nash equilibrium requires stringent conditions, posing a major challenge for this problem. To circumvent this difficulty, we examine the problem of identifying a mixed Nash equilibrium, where strategies are randomized and characterized by probability distributions over continuous domains. To this end, we propose PArticle-based Primal-dual ALgorithm (PAPAL) tailored for a weakly entropy-regularized min-max optimization over probability distributions. This algorithm employs the stochastic movements of particles to represent the updates of random strategies for the ϵ -mixed Nash equilibrium. We offer a comprehensive convergence analysis of the proposed algorithm, demonstrating its effectiveness. In contrast to prior research that attempted to update particle importance without movements, PAPAL is the first implementable particle-based algorithm accompanied by non-asymptotic quantitative convergence results, running time, and sample complexity guarantees. Our framework contributes novel insights into the particle-based algorithms for continuous min-max optimization in the general non-convex non-concave setting.

Keywords: Mixed Nash equilibrium, Particle-based algorithm.

*. Equal Contribution. Alphabetical order.

†. Work done at the Hong Kong University of Science and Technology.

1. Introduction

The problem of finding equilibrium of min-max optimization is a fundamental research topic, which has been a focus of numerous areas, such as mathematics, statistics, machine learning, economics, and computer science (Başar and Olsder, 1998; Roughgarden, 2010; Von Neumann and Morgenstern, 1947; Sinha et al., 2017). In machine learning, the applications of minimax optimization have shown great potential in real practice recently, such as generative adversarial networks (Goodfellow et al., 2014; Salimans et al., 2016), adversarial training (Ganin et al., 2016; Madry et al., 2017), reinforcement learning (Busoniu et al., 2008; Silver et al., 2017; Omidshafiei et al., 2017). These applications reveal the potential of broader learning systems beyond empirical risk minimization regime (Vapnik, 1991). Unfortunately, the theoretical properties of these objectives are not well studied, since min-max problems in real problems are usually non-convex non-concave. Practitioners often resort to heuristically implementing gradient-based algorithms for convex-concave min-max problems, which are convenient and scalable. Despite the potential and popularity of the applications with min-max optimization, the theoretical analysis lags far behind: the investigation on how to find a global solution with a provable algorithm in general min-max problems remains preliminary. The inherent complexities of min-max optimization impede our understanding of certain unanticipated phenomena and problematic scenarios in practical applications, such as instability, divergence, and oscillation (Goodfellow et al., 2014; Salimans et al., 2016; Daskalakis et al., 2017; Mescheder et al., 2018; Daskalakis and Panageas, 2018). It is urgent to study these problems and propose provable algorithms.

The min-max problems are usually written as:

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} f(x, y), \quad (1)$$

where $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ is the payoff function on the strategy space $\mathcal{X} \times \mathcal{Y}$. The solution to Eq. (1), termed the pure Nash equilibrium, is a strategy pair (x_*, y_*) that both min and max side cannot improve the corresponding loss value unilaterally. The Nash equilibrium is a single point pair in the strategy space. However, it exists only under very restrictive conditions (Dasgupta and Maskin, 1986). For general min-max problems, such as non-convex non-concave case, the pure Nash equilibrium does not usually exist (Arora et al., 2017). Given these challenges, researchers have explored alternative solutions to general min-max problems. One potential direction is to solve the general min-max problem in the sense of local equilibrium under varied assumptions (Evtushenko, 1974; Daskalakis and Panageas, 2018; Adolphs et al., 2019; Fiez et al., 2019; Mazumdar et al., 2020; Jin et al., 2020). These assumptions are often strong. For example, Fiez et al. (2019) propose “Differential Stackelberg Equilibrium” which restricts the behavior of the Hessian matrix. Diakonikolas et al. (2021) assume the minty variational inequality for the local Nash equilibrium, which is analogous to the Polyak-Lojasiewicz inequality in min-max problems. Mangoubi and Vishnoi (2021) consider the local case with several regularity conditions that the function f is smooth, bounded with Lipschitz Hessian. It was shown by Daskalakis et al. (2021) that finding a local min-max Nash equilibrium without these restrictions is also computationally expensive in general. Moreover, a local equilibrium is more or less limited compared with a global one.

Another criterion is called *mixed Nash equilibrium* (MNE), which proposes a random strategy over the space. Consider the extension of Eq. (1) as the following problem,

$$\min_{p \in \mathcal{P}_2(\mathcal{X})} \max_{q \in \mathcal{P}_2(\mathcal{Y})} \mathbb{E}_{p(x)} \mathbb{E}_{q(y)} [f(x, y)] + \mathcal{R}(p) - \mathcal{R}(q), \quad (2)$$

where the optimization variables p and q are densities of probability measures over \mathcal{X} and \mathcal{Y} , respectively, and $\mathcal{R}(\cdot)$ is a proper regularizer (such as relative entropy) to make the distributions well-behaved. By choosing some prior distribution π , the entropy regularizer can be written as

$$\mathcal{R}(p) = \lambda \mathbb{E}_p \log p - \lambda \mathbb{E}_p \log \pi, \quad (3)$$

where $\lambda > 0$ is the parameter to control the strength of the regularization term. The goal of mixed Nash equilibrium (MNE) is to find a probability distribution pair (p_*, q_*) in the strategy space that balances both sides. The regularization term ensures the problem is mathematically well-posed, and its influence becomes negligible as $\lambda \rightarrow 0$. Compared with the pure Nash equilibrium, the existence of MNE in min-max problems is guaranteed in the infinite-dimensional compact space (Glicksberg, 1952).

Finding the MNE with a provable algorithm is highly non-trivial since Eq. (2) is an infinite-dimensional problem. There have been some early attempts to solve the problem, where the core idea is to use a batch of particles (random samples) to parameterize the strategy and optimize the particles to approximate a MNE. Balandat et al. (2016) extend dual averaging to the min-max problem in continuous probability space, but the convergence rate is not discussed. Hsieh et al. (2019) propose a mirror-descent algorithm on probability measures to find MNEs. They show that a particle-based algorithm performs well in min-max problems empirically. However, the theoretical results for their algorithm remain open. Domingo-Enrich et al. (2020) attempt to propose a mean-field algorithm that uses both the weights and positions of the particles to parameterize the mixed strategy. However, in their analysis, particles do not move nearly as the weights' updates dominate the particles' movements in their setting. Such an algorithm will require a large number of particles, which is not implementable for high-dimensional problems.

Recent advances give us inspiration to seek the ϵ -mixed Nash equilibrium (ϵ -MNE) of Eq. (2), Nitanda et al. (2021) proposed a particle dual average algorithm with convergence guarantee, which optimizes the two-layer network with the second-order moment and negative entropy regularization successfully. Some other works focus on the fast convergence of continuous mean-field GDA dynamics (Lu, 2022), the convergence of quasistatic Wasserstein gradient flow (Ma and Ying, 2022) and the local convergence of particle algorithm (Wang and Chizat, 2022). Therefore, we aim to address a research question:

Can we design optimization algorithm that provides provable, quantitative guarantees for solving ϵ -MNE in continuous probability space?

In this paper, we first prove the existence of ϵ -MNE with relative entropy regularizers. We also propose the PArTicle-based Primal-dual ALgorithm (PAPAL), a mirror proximal primal-dual algorithm in the continuous probability space, where the KL divergence is chosen as the corresponding Bregman divergence (wrt the primal-dual algorithm). Unlike the finite

dimensional case, the sub-problems from the proximal operator in continuous probability space cannot be obtained in closed form. As a result, PAPAL is a double-looped algorithm, where the outer loop performs a mirror proximal primal-dual algorithm and the loop is approximated by the particle solution of the sub-problems. It is important to note that the challenges of particle-based algorithms primarily stem from errors introduced by both the finite-particle parameterization and the inexact solver.

The early works (Balandat et al., 2016; Hsieh et al., 2019) provide the idea of mirror descent in continuous probability space, but they only consider the continuous dynamics on the probability space. More recently, Domingo-Enrich et al. (2020) construct a dynamics of particles driven by both particle updates and reweighting. However, the underlying dynamics is dominated by weight updates, which means the proposed algorithm nearly does NOT move particles. Optimization within the probabilistic space, from their perspective, employs importance sampling. This approach is recognized for its vulnerability to the curse of dimensionality (Donoho et al., 2000). Thus, it is still not a real “implementable” algorithm with theoretical guarantees. In essence, both the inexact sub-problem solution and the functional estimation by finite particles are not discussed in previous work, which capture the behavior of computation complexity and sample size in real problems. To fill the gap, we show that such errors can be controlled to guarantee global convergence of the primal-dual algorithm when the prior distribution admits $\pi \propto \exp(-\lambda' \|x\|^2)$ (Gaussian distribution), where $\lambda' > 0$ controls the variance of the prior.

We also highlight the connection between PAPAL and PDA (Nitanda et al., 2021), in terms of their high-level motivation. Initially developed for constructing sequences within probability space and demonstrating convergence properties for mean-field problems with polynomial sample complexity, PDA served as a source of inspiration for our research. PAPAL surpasses the scope of PDA by extending its applicability to general min-max problems. From the perspective of convergence rate, without inner-loop error, PAPAL can achieve linear convergence rate. This relies on the momentum structure in the algorithm and sophisticated selection of hyperparameters. However, the convergence rate of PDA is significantly slower than linear rate. Specifically, PAPAL introduces a novel sequence construction methodology that diverges from the conventional emphasis on mean-field problems. We aim to extend the concept by devising multiple sub-problems in minimax optimizations. To achieve this, we have made modifications to the iterative formulation of the primal-dual approach to tailor it to particle approximation, while introducing novel analytical techniques beyond the realm of optimality conditions. To the best of our knowledge, we present the first particle-based algorithm that is not only practically implementable but also capable of approximating an MNE in continuous min-max optimization problems. Moreover, we provide a rigorous global convergence guarantee for our proposed approach.

Contributions. We propose PAPAL to solve the general ϵ -mixed Nash equilibrium (ϵ -MNE) in two-player zero-sum continuous games with minimal regularization. We validate the existence of MNE within our framework. Our analysis focuses on the convergence of the primal-dual method in addressing MNE. Specifically, we demonstrate that the sampling sub-problems exhibit log-Sobolev properties, and we can ensure global convergence even when utilizing inexact solvers for sub-problems. We present quantitative convergence metrics, runtime estimations, and sample complexity for PAPAL. Additionally, we elucidate the

analysis of the stochastic gradient variant of the PAPAL algorithm. To our knowledge, this is the first demonstration of the computational and sample feasibility of a primal-dual method in a continuous probability framework to resolve min-max challenges.

2. Related Work

Min-Max Optimization. The study of solving min-max problems is quite extensive since its inception with (Von Neumann, 1928). The gradient-based algorithms, including those presented by Korpelevich (1976); Nemirovski and Yudin (1978); Chen and Rockafellar (1997); Nemirovski (2004), leverages convexity and concavity to obtain the convergence properties. Notably, gradient descent ascent (GDA) finds an ϵ -approximate stationary point within $\mathcal{O}(\log 1/\epsilon)$ iterations for strongly-convex strongly-concave minimax problems, and $\mathcal{O}(\epsilon^{-2})$ iterations for convex concave games (Nemirovski, 2004; Nedić and Ozdaglar, 2009). While some alternative gradient methodologies extend to non-convex concave scenarios (Namkoong and Duchi, 2016; Grnarova et al., 2017; Rafique et al., 2019; Lu et al., 2020; Lin et al., 2020), the general min-max domain remains relatively uncharted. A central challenge emerges from the dependency of general min-max problems on stringent conditions, coupled with a lack of consensus on assumptions. By worst-case evaluations, non-convex non-concave min-max problems can be computationally formidable. Daskalakis et al. (2009) highlighted the PPAD-completeness of establishing a general pure Nash equilibrium without the condition of convexity/concavity. Similarly, Chen et al. (2006) prove the PPAD-completeness for non-zero two-player mixed Nash equilibria in bimatrix game, and subsequent studies (Daskalakis, 2013; Rubinstein, 2016) underlined this complexity even for approximated mixed Nash equilibria. Finding an (ϵ, δ) -local min-max Nash equilibrium in non-convex non-concave min-max problems is also PPAD-complete (Daskalakis et al., 2021). While recent endeavors (Daskalakis and Panageas, 2018; Adolphs et al., 2019; Mazumdar et al., 2020; Jin et al., 2020; Mangoubi and Vishnoi, 2021) have delved into the properties of local min-max points in non-convex frameworks, real-world interpretations of the imposed conditions remain elusive. Consequently, general pure Nash Equilibrium’s feasibility is widely regarded as remote, barring justifiable restrictive conditions.

Mixed Nash Equilibrium. Mixed Nash Equilibrium lifts the strategy space to a probability space and convert min-max problem to be infinite-dimensional. Importantly, the existence of Mixed Nash Equilibrium can be guaranteed when the payoff function is differentiable with Lipschitz continuous gradient (Glicksberg, 1952). Recently, with the growing importance of general min-max problems in deep learning and Generative Adversarial Networks (GAN), the investigation of Mixed Nash Equilibrium has gained more attention due to its good theoretical properties. The global existence without convexity assumption make it more reasonable in real problems. Grnarova et al. (2017) study the behavior in a special setting: they consider the online setting of GANs with a two-layer neural network as the discriminator. Arora et al. (2017) study the generalization properties of the Mixed Nash Equilibrium in GAN. Moreover, there is a line of works that try to solve the Mixed Nash Equilibrium with mirror-descent-like algorithms in continuous probability space. Balandat et al. (2016) first solve Mixed Nash Equilibrium with dual averaging algorithm, but the convergence rate is not discussed. Earlier discussions between (Hsieh et al., 2019) and (Domingo-Enrich et al., 2020) render repetition unnecessary. In a word, there are some

early attempts to perform mirror proximal algorithms in probability space. In essence, while preliminary endeavors into mirror proximal algorithms in the probability space have emerged, with empirical evidence underscoring their significance, a rigorous theoretical dissection of "implementable" particle-based algorithms is still in its nascence.

Mean-field Analysis. Mean field analysis focuses on the evolution of particles on infinite-dimensional space (Dobrushin, 1979), which can be described by some non-linear partial differential equations (PDE). Notably, many non-convex dynamical systems in modern machine learning can be analyzed by lifting the original problem to infinite-dimensional space. Such a procedure convexifies the original problem, which has better mathematical properties (Bengio et al., 2005; Bach, 2017). The limiting PDE in mean-field analysis can derive the global convergence and well-posed behaviors. For example, the mean-field analysis over the neurons in wide neural networks captures the feature evolution of the optimization process of deep learning (Mei et al., 2018, 2019; Yang and Hu, 2020; Nguyen and Pham, 2020; Fang et al., 2021; Yang et al., 2022; Fang et al., 2022). There have been a number of works based on noisy gradient descent to investigate the quantitative convergence rate when minimizing the nonlinear functionals with entropy regularizers (Hu et al., 2021; Nitanda et al., 2021; Chizat, 2022). In summary, mean-field analysis stands as a powerful theoretical tool in today's machine learning theory, which helps to convexify high-dimensional problems through randomization (a special lifting to infinite dimensional space).

3. Preliminary

Notation. We use lower-case letters x, y to denote vectors. For function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, $\nabla f(\cdot)$ and $\nabla^2 f(\cdot)$ denote its gradient and Hessian matrix, respectively. For function $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$, $\nabla f(\cdot)$ denotes the Jacobian matrix. For multivariate function $f(x, y)$, ∇_x, ∇_y denote the gradient over x, y . Notation $\|\cdot\|$ denotes 2-norm for both vector and matrix. The continuous probability measures in this paper are considered to be absolutely continuous by default with respect to the Lebesgue measure, which induces density function p and q . The smoothness is guaranteed by Gaussian initialization and Gaussian smoothing during the updates. For twice differentiable function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we say that f is L_0 -Lipschitz if $|f(x) - f(y)| \leq L_0 \|x - y\|$ for any $x, y \in \mathbb{R}^d$; we say that f has L_1 -Lipschitz gradient if $\|\nabla f(x) - \nabla f(y)\| \leq L_1 \|x - y\|$ for any $x, y \in \mathbb{R}^d$. For any a topological space \mathcal{X} , $\mathcal{P}(\mathcal{X})$ denotes set of probability measures on \mathcal{X} and $\mathcal{P}_2(\mathcal{X})$ denotes the set of probability measures with bounded second moments. For density p , we define $\mathcal{E}(p) = \mathbb{E}_p[\log(p)]$, $\text{KL}(p||q) = \mathbb{E}_p[\log(p/q)]$. For the generalized density for particles, we use Dirac delta distribution δ_x to represent the Dirac delta distribution at x , $\delta_{\{x_1, \dots, x_n\}} = n^{-1} \sum_{i=1}^n \delta_{x_i}$.

Problem Setup. Assume that the min-max problem is defined on the joint parameter space $\mathcal{X} \times \mathcal{Y}$ and a payoff function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$. We consider the bounded function case. For simplicity without loss of generality, the domain $\mathcal{X} \times \mathcal{Y}$ is the Euclidean space $\mathbb{R}^m \times \mathbb{R}^n$ and the co-domain of f is a bounded interval $[-1, 1]$.

In our work, we consider the regularized minimax problem defined in Eq. (2). For the regularization term, we set

$$\mathcal{R}(p(x)) = \mathbb{E}_{p(x)} [\lambda_1 \|x\|^2 + \lambda_2 \log(p(x))] \quad (4)$$

where $\lambda_1, \lambda_2 > 0$ are constants to determine the regularization scaling as well as the variance of the prior. We use $\mathcal{L}(p, q)$ to denote $\mathbb{E}_{p(x)}\mathbb{E}_{q(y)}f(x, y) + \mathcal{R}(p) - \mathcal{R}(q)$. The mixed Nash equilibrium denotes a distribution pair (p_*, q_*) such that for all $(p, q) \in \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y})$, $\mathcal{L}(p_*, q) \leq \mathcal{L}(p_*, q_*) \leq \mathcal{L}(p, q_*)$. Therefore, for any distribution pair (\bar{p}, \bar{q}) , the duality gap $\max_{q \in \mathcal{P}(\mathcal{Y})} \mathcal{L}(\bar{p}, q) - \min_{p \in \mathcal{P}(\mathcal{X})} \mathcal{L}(p, \bar{q})$ forms the basis for a standard optimality criterion. Formally, we define ϵ -MNE as below.

Definition 1 (ϵ -MNE). *A distribution pair (\bar{p}, \bar{q}) is an ϵ -mixed Nash equilibrium (ϵ -MNE) of function $\mathcal{L}(\cdot, \cdot)$ if*

$$\max_{q \in \mathcal{P}(\mathcal{Y})} \mathcal{L}(\bar{p}, q) - \min_{p \in \mathcal{P}(\mathcal{X})} \mathcal{L}(p, \bar{q}) \leq \epsilon.$$

When $\epsilon = 0$, (\bar{p}, \bar{q}) is a mixed Nash equilibrium.

Assumptions.

[A₁] There exists L_0 such that f is L_0 -Lipschitz.

[A₂] f has L_1 -Lipschitz gradient.

[A₃] For any $x, y \in \mathcal{X} \times \mathcal{Y}$, $|f(x, y)| \leq 1$.

We consider the bounded function with growth regularities. The bound of $|f|$ is chosen as 1 for notation simplicity. Assumptions **[A₁]** and **[A₂]** regularize the growth of f , ∇f respectively and the definition of Lipschitz continuity was introduced in **Notation** part. It is clear that our assumption is quite weak so our setting is general.

Primal-Dual Algorithms for Min-Max Optimization. Primal-dual method is a classical approach to solve saddle point problems (Nemirovski, 2004; He et al., 2017; Lan, 2020; Bredies and Sun, 2015; Luke and Shefi, 2018). Consider the min-max problem

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} \{f(x) + \langle Ax, y \rangle - g(y)\}, \quad (5)$$

where $A \in \mathbb{R}^{m \times n}$ denotes a linear mapping, and $f : \mathcal{X} \rightarrow \mathbb{R}$ and $g : \mathcal{Y} \rightarrow \mathbb{R}$ are convex functions. General optimization problems with linear constraints: $\min_{x \in \mathcal{X}} \{f(x) : Ax = b\}$ can be transformed to Eq. (5). The proximal primal-dual algorithm updates x_t and y_t by

$$\tilde{x}_t = x_{t-1} + \mu_t(x_{t-1} - x_{t-2}), \quad (6)$$

$$y_t = \arg \min_{y \in \mathcal{Y}} \langle -A\tilde{x}_t, y \rangle + g(y) + \tau_t W(y_{t-1}, y), \quad (7)$$

$$x_t = \arg \min_{x \in \mathcal{X}} \langle y_t, Ax \rangle + f(x) + \eta_t V(x_{t-1}, x), \quad (8)$$

under a finite dimension condition with W and V representing Bregman's distances on \mathcal{Y} and \mathcal{X} , respectively. Moreover, τ_t, η_t, γ_t are algorithmic parameters to be specified. The Bregman divergence is particularly important for other Riemannian structures. In probability space, the Bregman divergence is usually chosen as KL divergence or Wasserstein distance.

Unadjusted Langevin Algorithm. We introduce the *Unadjusted Langevin Algorithm (ULA)* to solve the mirror steps in our setting. In particular, some isoperimetry inequalities shall be introduced to make each step well-posed.

Definition 2. (*Gross, 1975*) We say that density p satisfies *log-Sobolev inequality (LSI)* with a constant $\alpha > 0$, if

$$\mathbb{E}_p[g^2 \log(g^2)] - \mathbb{E}_p[g^2] \log \mathbb{E}_p[g^2] \leq \frac{2}{\alpha} \mathbb{E}_p[\|\nabla g\|^2] \quad (9)$$

for all differentiable function g with $\mathbb{E}_p[g^2] < \infty$. When $g = \sqrt{q/p}$, we have

$$\text{KL}(q\|p) \leq \frac{1}{2\alpha} \mathbb{E}_q \left\| \nabla \log \frac{q}{p} \right\|^2.$$

Assume we aim to sample from a smooth target probability density $p \propto e^{-f(x)}$ where $f(x)$ has L_1 -Lipschitz gradient and satisfies a smoothness condition. The ULA with step size $\epsilon > 0$ is the discrete-time algorithm

$$x_{s+1} = x_s - \epsilon \nabla f(x_s) + \sqrt{2\epsilon} \zeta_s, \quad (10)$$

where $s = 0, \dots, T$ for some positive integer T ; x_0 is given as an initialization; $\zeta_s \sim \mathcal{N}(0, I)$ is an isotropic Gaussian noise. Moreover, if p satisfies LSI, we can guarantee the convergence rate of ULA. The $\text{KL}(\rho_s\|p)$ along ULA would also converge to $\mathcal{O}(\epsilon)$ with $\mathcal{O}(\epsilon)$ fixed step size. Since it is equivalent with gradient descent with additional Gaussian noise, ULA is easy to implement in practice. The easy-to-implement property and the arbitrarily small bias make it a popular in sampling literature. Thus, we choose it to solve our sub-problems.

4. PAPAL: Particle-based Primal-dual ALgorithm

We propose Particle-based Primal-dual ALgorithm (PAPAL) to solve the regularized minimax problem Eq. (2) by applying Mirror-Prox Primal-dual method to MNE.

As an optimization algorithm on the space of continuous probability measures, our proposed PAPAL (Algorithm 1) consists of two components: inner loop and outer loop. Specifically, our outer loop extends the classical primal-dual scheme (Sec. 3) to infinite dimensional optimization problems (Sec. 4.1). Each outer step can be rewritten as a sampling task for Gibbs distribution (Sec. 4.2). We perform ULA as the inner loop to approximate a proximal Gibbs distribution, which is applied to the outer loop optimization to converge to the Mixed Nash Equilibrium distributions p^* and q^* (Sec. 4.3). We will provide the motivation of our algorithm from an optimization view (Sec. 4.4).

4.1 Exact Primal-dual Method

We construct the target density of p_t^* and q_t^* for each iteration. Particularly, given \hat{p}_{-1} , \hat{p}_0 , and \hat{q}_0 , $q_t^*(y)$ and $p_t^*(x)$ satisfy the following recursive condition,

$$q_t^* := \operatorname{argmin}_{q \in \mathcal{P}_2(\mathcal{Y})} \mathbb{E}_q[\hat{\phi}_t(y)] + \mathcal{R}(q) + \tau_t \text{KL}(q\|q_{t-1}^*), \quad (11)$$

$$p_t^* := \operatorname{argmin}_{p \in \mathcal{P}_2(\mathcal{X})} \mathbb{E}_p[\hat{\psi}_t(x)] + \mathcal{R}(p) + \eta_t \text{KL}(p\|p_{t-1}^*), \quad (12)$$

Algorithm 1 PAPAL: PArticle-based Primal-dual ALgorithm for Mixed Nash Equilibrium

Input: number of outer iterations T , outer-loop learning rates $\{\eta_t\}_{t=1}^T$ and $\{\tau_t\}_{t=1}^T$, outer-loop hyper-parameters $\{\mu_t\}_{t=1}^T$; number of inner iterations $\{T_t\}_{t=1}^T$, inner-loop learning rates $\{\iota_t\}_{t=1}^T$, number of particles M

Output: estimated equilibrium $(\frac{1}{M}\delta_{\hat{X}}, \frac{1}{M}\delta_{\hat{Y}})$

- 1: Initialize $\{\tilde{x}_r^{(0)}\}_{r=0}^M \sim p_0(x) \propto \exp\left\{-\frac{\lambda_1}{\lambda_2}\|x\|^2\right\}$, $\{\tilde{y}_r^{(0)}\}_{r=0}^M \sim q_0(y) \propto \exp\left\{-\frac{\lambda_1}{\lambda_2}\|y\|^2\right\}$.
 Assign $(\tilde{X}^{(-1)}, \tilde{Y}^{(-1)}) = (\tilde{X}^{(0)}, \tilde{Y}^{(0)}) \leftarrow (\{\tilde{x}_r^{(0)}\}_{r=1}^M, \{\tilde{y}_r^{(0)}\}_{r=1}^M)$.
 Assign $h^{(0)}(y) = \frac{\lambda_1}{\lambda_2}\|y\|^2$, $g^{(0)}(x) = \frac{\lambda_1}{\lambda_2}\|x\|^2$.
 - 2: **for** $t = 0$ to $T - 1$ **do**
 - 3: Initialize $(X^{(0)}, Y^{(0)}) \sim p_0(x) \times q_0(y)$.
 - 4: Update $h^{(t+1)}(y) = \frac{1}{\lambda_2 + \tau_{t+1}} \left\{ \hat{\phi}_{t+1}(y) + \lambda_1\|y\|^2 + \tau_{t+1}h^{(t)}(y) \right\}$.
 - 5: **for** $k = 0$ to T_t , $r = 0$ to M **do**
 - 6: Run noisy GD $y_r^{(k+1)} \leftarrow y_r^{(k)} - \iota_t \nabla_y h^{(t+1)}(y_r^{(k)}) + \sqrt{2\iota_t} \xi_r^{(k)}$, where $\xi_r^{(k)} \sim \mathcal{N}(0, I_n)$.
 - 7: **end for**
 - 8: Assign $\tilde{Y}^{(t+1)} \leftarrow Y^{(T_t+1)}$.
 - 9: Update $g^{(t+1)}(x) = \frac{1}{\lambda_2 + \eta_{t+1}} \left\{ \hat{\psi}_{t+1}(x) + \lambda_1\|x\|^2 + \eta_{t+1}g^{(t)}(x) \right\}$.
 - 10: **for** $k = 0$ to T_t , $r = 0$ to M **do**
 - 11: Run noisy GD $x_r^{(k+1)} \leftarrow x_r^{(k)} - \iota_t \nabla_x g^{(t+1)}(x_r^{(k)}) + \sqrt{2\iota_t} \zeta_r^{(k)}$, where $\zeta_r^{(k)} \sim \mathcal{N}(0, I_m)$.
 - 12: **end for**
 - 13: Assign $\tilde{X}^{(t+1)} \leftarrow X^{(T_t+1)}$.
 - 14: **end for**
 - 15: **Option I:** $(\frac{1}{M}\delta_{\hat{X}}, \frac{1}{M}\delta_{\hat{Y}}) \leftarrow (\frac{1}{M}\delta_{\tilde{X}^{(T)}}, \frac{1}{M}\delta_{\tilde{Y}^{(T)}})$.
 - 16: **Option II:** Randomly pick up $(\frac{1}{M}\delta_{\tilde{X}^{(t)}}, \frac{1}{M}\delta_{\tilde{Y}^{(t)}})$ from $t \in \{1, \dots, T\}$ following the probability $\mathbb{P}[t] = \frac{\mu_t}{\sum_{t=1}^T \mu_t}$ as $(\frac{1}{M}\delta_{\hat{X}}, \frac{1}{M}\delta_{\hat{Y}})$.
-

where $\hat{\phi}_t(y) = -(1 + \mu_t)\mathbb{E}_{\hat{p}_{t-1}}[f(\cdot, y)] + \mu_t\mathbb{E}_{\hat{p}_{t-2}}[f(\cdot, y)]$, and $\hat{\psi}_t(x) = \mathbb{E}_{\hat{q}_t}[f(x, \cdot)]$.

If $\hat{q}_t = q_t^*$, $\hat{p}_t = p_t^*$, Eq. (11) and (12) perform exact primal-dual method in continuous probability space, with KL divergence as the Bregman divergence. It is quite direct to compare them to Eq. (7) and (8). However, there two gaps for the aforementioned updates: (1) the forms of q_{t-1}^* and p_{t-1}^* are not explicit in Eq. (11) and Eq. (12); (2) the exact solutions for $\hat{q}_t \equiv q_t^*$ and $\hat{p}_t \equiv p_t^*$ are not tractable. Thus, it is necessary to derive a explicit form of (q_{t-1}^*, p_{t-1}^*) and approximate the optimal (q_t^*, p_t^*) with finite M particles in (\hat{q}_t, \hat{p}_t) .

4.2 Recursive Energy Function

Section 4.1 provides us an algorithm to find mixed Nash equilibrium in the probability space. In this part, we would discuss the feasibility of the proposed updates.

The following Lemma provides us insight to obtain the explicit form of q_*^{t-1} and p_*^{t-1} in Eq. (11) and Eq. (12) recursively, which can be implemented in an algorithm.

Lemma 3. *Let $\lambda > 0$ be a positive real number and $\tilde{l}(\theta), l(\theta)$ be bounded continuous functions. Consider a probability density $\bar{p}(\theta) \propto \exp(-\tilde{l}(\theta))$, then $p(\theta) \propto \exp\left(-\frac{1}{\lambda+\tau}l(\theta) - \frac{\tau}{\lambda+\tau}\tilde{l}(\theta)\right)$ is an optimal solution to the following problem*

$$\min_{p \in \mathcal{P}_2(\mathbb{R}^{d_\theta})} \{ \mathbb{E}_p[l(\theta)] + \lambda \mathcal{E}(p) + \tau \text{KL}(p \| \bar{p}) \}. \quad (13)$$

Once we obtained two approximate distributions $\hat{p}_i(y)$ and $\hat{p}_i(x)$, Lemma 3 indicates that if p_{t-1}^* and q_{t-1}^* are proximal Gibbs distribution with explicit form. Then, p_t^* and q_t^* have the similar expression as well. Therefore, we denote

$$h^{(t)}(y) = \frac{1}{\lambda_2 + \tau_t} \left\{ \hat{\phi}_t(y) + \lambda_1 \|y\|^2 + \tau_t h^{(t-1)}(y) \right\}, \quad (14)$$

$$g^{(t)}(x) = \frac{1}{\lambda_2 + \eta_t} \left\{ \hat{\psi}_t(x) + \lambda_1 \|x\|^2 + \eta_t g^{(t-1)}(x) \right\}, \quad (15)$$

which consist of the Gibbs potential on $\mathcal{P}_2(\mathcal{Y})$ and $\mathcal{P}_2(\mathcal{X})$ after $t - 1$ iterations. Thus, at each step- t , our objective probability densities can be written as

$$q_t^*(y) = \frac{\exp(-h^{(t)}(y))}{Z_{t,1}}, \quad p_t^*(x) = \frac{\exp(-g^{(t)}(x))}{Z_{t,2}},$$

where $Z_{t,1}$ and $Z_{t,2}$ are normalizing constants.

4.3 Inexact Solver and Particle Approximation

At every step t , achieving the ideal scenario where $\hat{q}_t = q_t^*$ and $\hat{p}_t = p_t^*$ is not tractable within the algorithm. This limitation arises because obtaining exact solutions for Equations (11) and (12) is not feasible for a general function f . Fortunately, the optimization involving random particles can be effectively approximated using the log density. ULA presents a viable approach for approximating the desired target distribution. By performing ULA, we can obtain random particle samples $(\{x_r^{T_t}\}_{r=1}^M, \{y_r^{T_t}\}_{r=1}^M)$ that can approximate (q_t^*, p_t^*) . We let that $(p_t, q_t) = (\text{Law}(\{x_r^{T_t}\}_{r=1}^M), \text{Law}(\{y_r^{T_t}\}_{r=1}^M))$ and (\hat{q}_t, \hat{p}_t) is finite sample approximation of (p_t, q_t) (delta distribution by finite samples).

We first initialize a set of particles and duplicate it by setting $\{\tilde{x}_r^{(-1)}, \tilde{y}_r^{(-1)}\}_{r=1}^M = \{\tilde{x}_r^{(0)}, \tilde{y}_r^{(0)}\}_{r=1}^M$. For each step, we run ULA to optimize the particles in \mathcal{Y} , where the gradient at the k -th inner step is given by $\nabla_y h^{(t)}(y_r^{(k)})$ estimated by $\{\tilde{x}_r^{(t-1)}\}_{r=1}^M$ and $\{\tilde{x}_r^{(t-2)}\}_{r=1}^M$. Then, we execute the same process as above starting from the prior where the gradient at the k -th inner step is given by $\nabla_y g^{(t)}(x_r^{(k)})$ estimated by $\{\tilde{y}_r^{(t)}\}_{r=1}^M$. Thus, we obtain the full procedure of PAPAL in Algorithm 1. We also provide empirical results in Appendix G.

Remark 4. *To make a clear distinction between (\hat{q}_t, \hat{p}_t) , (q_t, p_t) , and (q_t^*, p_t^*) , we reiterate the definition here: we initialize $(\hat{q}_{-1}, \hat{p}_{-1}) = (\hat{q}_0, \hat{p}_0)$ from the prior distribution. From $t = 1$ to $t = T$, (q_t^*, p_t^*) is the exact solution from Eq. (11) and Eq. (12); (q_t, p_t) is the ULA approximation of (q_t^*, p_t^*) ; (\hat{q}_t, \hat{p}_t) is the realization of (q_t, p_t) with M samples (mixture of dirac delta distributions). All the densities can be defined by recursion.*

Remark 5. *Other sampling algorithms can be implemented as the inexact solver, such as Hamiltonian Monte Carlo (Duane et al., 1987; Neal et al., 2011), Metropolis-adjusted Langevin algorithm (Rosicky et al., 1978; Grenander and Miller, 1994). We use ULA here for demonstration and simplicity. More sophisticated solvers may accelerate the convergence.*

4.4 Optimization View of PAPAL

In this subsection we explain the intuition behind PAPAL from an optimization perspective. Recall that Algorithm 1 iteratively updates the particles \hat{q}_t and \hat{p}_t which approximate

the minimizers of the linearized potentials Eq. (11) and (12). Actually, we obtain the classical Primal-dual algorithm (Nemirovski, 2004) if Eq. (11) and Eq. (12) degenerate to finite dimension. It is obvious that we can find q_t^* and p_t^* with negligible error under finite dimension condition which implies the total complexity would be derived easily. However, there are three problems if we rely on Langevin algorithm when solving the sub-problem for infinite dimension:

1. Express the objective densities of the sub-problem as the form of $e^{-h^{(t)}(y)}/Z_{t,1}$ and $e^{-g^{(t)}(x)}/Z_{t,2}$, which can be explicitly obtained in an algorithm;
2. Analysis the outer loop with the bounded KL error induced by ULA at each step- t , which converges linearly with fixed step size under LSI;
3. Control the error induced by finite particles at each step- t , which should be done by concentration.

The first problem can be solved by forward recursion of (q_{t-1}^*, p_{t-1}^*) and Lemma 3 combining the approximated solutions $\{(\hat{q}_k, \hat{p}_k)\}_{k=-1}^{t-1}$ we have known. Next, in order to bound the error between approximate densities and precise densities in each iteration of Primal-dual method, we introduce the LSI condition (Gross, 1975), which can guarantee the fast convergence of sampling algorithms, such as ULA. Therefore, we add $\|\cdot\|^2$ term with coefficient λ_1 , (λ_1 can be an arbitrary positive number) to the regularizer which guarantee that q_t^* and p_t^* satisfy the LSI condition. Finally, the finite particle $\hat{\phi}_t$ and $\hat{\psi}_t$ approximate the true expectation $\phi_t(y) = -(1 + \mu_t)\mathbb{E}_{p_{t-1}}[f(\cdot, y)] + \mu_t\mathbb{E}_{p_{t-2}}[f(\cdot, y)]$, and $\psi_t(x) = \mathbb{E}_{q_t}[f(x, \cdot)]$. The gap between $(\hat{\phi}_t, \hat{\psi}_t)$ and (ϕ_t, ψ_t) shall be controlled by concentration techniques.

To sum up, we connect bilinear forms $\langle Ax, y \rangle$ under finite dimension and $\mathbb{E}_{p(x)}\mathbb{E}_{q(y)}f(x, y)$ under infinite dimension with proximal Gibbs distribution. In addition, we apply the estimation of sub-problem error and finite sample error to global convergence analysis for our particle-based algorithm. Therefore, we may interpret the PAPAL as a finite-particle approximation of a primal-dual method. PAPAL combines the efficient convergence of Primal-dual method in outer loop and the KL error bound of ULA under LSI condition in inner loop. More importantly, the loss functional can be estimated with polynomial finite particles, which makes our algorithm implementable.

5. Existence of MNE

Our objectives are two-fold: 1) To achieve an ϵ -approximate Mixed Nash Equilibrium (MNE) using the PAPAL algorithm. 2) To minimize the distance between the continuous probability density function generated by the algorithm and the MNE of the problem defined in Eq. (2), with the regularizer outlined in Eq. (4), aiming to reduce this gap as close to zero as possible. A fundamental step towards these goals is establishing the existence of the MNE for the problem in Eq. (2). To this end, the subsequent theorem asserts the existence of an MNE for this problem, considering the regularization term in Eq. (4). It is worth noting that the existence result in this paper is established in $\mathbb{R}^m \times \mathbb{R}^n$, while the existence result proposed by Domingo-Enrich et al. (2020) is based on compact Riemann manifolds without boundary embeded in $\mathbb{R}^m \times \mathbb{R}^n$.

Theorem 6. *Considering problem Eq. (2) with regularizer Eq. (4) and assuming Assumptions [A₁], [A₂] and [A₃] hold, there exists a MNE (p_*, q_*) satisfies that*

$$\begin{aligned} p_* &= \arg \min_{p \in \mathcal{P}_2(\mathcal{X})} \mathbb{E}_{p(x)} \mathbb{E}_{q_*(y)} [f(x, y)] + \mathcal{R}(p) - \mathcal{R}(q_*), \\ q_* &= \arg \min_{q \in \mathcal{P}_2(\mathcal{Y})} \mathbb{E}_{p_*(x)} \mathbb{E}_{q(y)} [f(x, y)] + \mathcal{R}(p_*) - \mathcal{R}(q). \end{aligned} \quad (16)$$

The proof of Theorem 6 can be formulated into a two-step process. Step 1: constructing a suitable probability density function family and leveraging well-established lemmas in functional analysis to verify that this function family is a compact convex set. Step 2: proving that the operator $\mathbf{F}(p, q)$ (as precisely defined in the Appendix B) is a continuous mapping from the aforementioned function family to itself. Finally, we can achieve the conclusion of Theorem 6 combining Brouwer fixed-point theorem.

6. Convergence Analysis

In this section, we provide quantitative global convergence for PAPAL. We first provide the finite sample error bound in Sec. 6.1. Then, under deterministic version, we introduce the outer loop settings in Sec. 6.2.1 and conclude the properties of the updates. By assuming the approximate optimality of the inner loop iterations, the error bound of the outer loop can be obtained. In Sec. 6.2.2, we discuss the convergence rate for the inner loop iterations, which is the KL convergence of ULA algorithm. In Sec. 6.2.3, we combine the outer loop and the inner loop to discuss the total complexity of deterministic PAPAL and provide the global convergence results of our main theorem. As a supplementary remark, Sec 6.3 presents a global convergence analysis of the stochastic PAPAL algorithm.

6.1 Finite Sample Error

We introduce two probability sets $\mathcal{Q}_{\mathcal{X}}(C)$ and $\mathcal{Q}_{\mathcal{Y}}(C)$ for positive constant C as follows

$$\begin{aligned} \mathcal{Q}_{\mathcal{X}}(C) &:= \left\{ p(x) \in \mathcal{P}_2(\mathcal{X}) \mid p(x) \propto e^{-l_1(x) - \frac{\lambda_1}{\lambda_2} \|x\|^2}; \|l_1\|_{\infty} \leq \frac{C}{\lambda_2} \right\}, \\ \mathcal{Q}_{\mathcal{Y}}(C) &:= \left\{ q(y) \in \mathcal{P}_2(\mathcal{Y}) \mid q(y) \propto e^{-l_2(y) - \frac{\lambda_1}{\lambda_2} \|y\|^2}; \|l_2\|_{\infty} \leq \frac{C}{\lambda_2} \right\}. \end{aligned}$$

To establish connections between ϕ_t and $\hat{\phi}_t$, as well as ψ_t and $\hat{\psi}_t$, in the convergence analysis, it is sufficient to bound the discretization error on linear term by using the form of objective function in Eq. (2). Therefore, we use the following lemma to estimate the discretization error between \hat{p}_t and p_t (\hat{q}_t and q_t) at each step- t (refer to Appendix F).

Lemma 7. *For any $\epsilon > 0, \delta \in (0, 1)$, the errors of finite particles satisfy*

$$\mathbb{P} [|(\hat{p}_t - p_t)(x) f(x, y) q(y) dy dx| \geq \epsilon] \leq \delta, \quad \mathbb{P} [|(\hat{q}_t - q_t)(y) f(x, y) p(x) dx dy| \geq \epsilon] \leq \delta, \quad (17)$$

with the required numbers of particles

$$M := \mathcal{O} \left(\epsilon^{-2} \left[-\log(\delta) + d \log \left(1 + L_0 r_{\max\{C_1, C_2\}}(\epsilon, \lambda_1, \lambda_2) \epsilon^{-1} \right) \right] \right)$$

on $\hat{p}_t(x)$ and $\hat{q}_t(y)$ for all $p(x) \in \mathcal{Q}_{\mathcal{X}}(C_1)$, $q(y) \in \mathcal{Q}_{\mathcal{Y}}(C_2)$ where $d = \max\{m, n\}$ and $r_C(\epsilon, \lambda_1, \lambda_2) := 8 \sqrt{\frac{\lambda_2 + C}{\lambda_1} \max\{\log(\epsilon^{-1}), d\}}$.

We introduce the max-min gap function

$$Q(\bar{w}, w) := \mathcal{L}(\bar{p}, q) - \mathcal{L}(p, \bar{q}), \quad (18)$$

where $\bar{w} := (\bar{p}, \bar{q})$ and $w := (p, q)$. Note that the maximum of Eq. (18) equals to the duality gap on (\bar{p}, \bar{q}) . Considering the objective function in Eq. (2) with regularizer (4), we have

$$\begin{aligned} \max_{w \in \mathcal{P}_2(\mathcal{X}) \times \mathcal{P}_2(\mathcal{Y})} Q(\bar{w}, w) &\geq \lambda_2 (\text{KL}(\bar{p} \parallel \bar{p}^*) + \text{KL}(\bar{q} \parallel \bar{q}^*)), \\ \text{s.t. } \bar{w}^* = (\bar{p}^*, \bar{q}^*) &= \operatorname{argmax}_{w' \in \mathcal{P}_2(\mathcal{X}) \times \mathcal{P}_2(\mathcal{Y})} Q(\bar{w}, w'). \end{aligned} \quad (19)$$

As a result, the max-min gap on \bar{w} can regulate the KL divergence between \bar{w} and \bar{w}^* .

To simplify the notation of the formula and capture the error of our algorithm, we use concise symbols to represent some complex intermediate variables: $r_C(\epsilon, \lambda_1, \lambda_2) := 8\sqrt{\frac{\lambda_2 + C}{\lambda_1} \max\{\log(\epsilon^{-1}), m, n\}}$ and $g(\lambda_2) := \log(1 + \lambda_2 + \lambda_2^2)$, and assume that the KL divergence between (q_t, p_t) and (q_t^*, p_t^*) can be bounded by small errors in outer loop, i.e.,

$$\text{KL}(p_t \parallel p_t^*) \leq \delta_{t,1}, \quad \text{KL}(q_t \parallel q_t^*) \leq \delta_{t,2},$$

which is guaranteed by ULA algorithm with small step sizes (refer to section 6.2.2). We use the error bound quantities $\delta_{t,1}$ and $\delta_{t,2}$ to analyze the outer loop properties.

6.2 Convergence of PAPAL

6.2.1 OUTER LOOP ERROR BOUND

The following theorem proves the convergence properties of our PAPAL.

Theorem 8. *By setting parameters $\gamma_t \mu_t = \gamma_{t-1}$, $\gamma_t \tau_t \leq \gamma_{t-1}(\tau_{t-1} + \lambda_2)$, $\gamma_t \eta_t \leq \gamma_{t-1}(\eta_{t-1} + \lambda_2)$, $\tau_t \eta_{t-1} \geq \mu_t$, $\mu_t \leq 1$ for any $t \geq 1$ and optimizing the sub-problem with errors bounded by $\delta_{t,1}, \delta_{t,2}$ mentioned above, iterates of our method satisfies*

$$\begin{aligned} \sum_{t=1}^T \gamma_t Q(w_t, w) &\leq \gamma_1 \eta_1 \text{KL}(p \parallel p_0^*) - \gamma_T (\eta_T + \lambda_2) \text{KL}(p \parallel p_T^*) \\ &\quad + \gamma_1 \tau_1 \text{KL}(q \parallel q_0^*) - \gamma_T \left(\tau_T + \lambda_2 - \frac{1}{4\eta_T} \right) \text{KL}(q \parallel q_T^*) \\ &\quad + 4 \sum_{t=1}^k \gamma_t (1 + \mu_t) \epsilon \\ &\quad + \sum_{t=1}^k \gamma_t \left[\left(\lambda_2 + \frac{2\lambda_1}{\alpha_{p_t^*}} \right) \delta_{t,1} + \left(30 + \frac{8\lambda_1 \sigma(p_t^*)}{\sqrt{\alpha_{p_t^*}}} \right) \sqrt{\delta_{t,1}} \right] \\ &\quad + \sum_{t=1}^k \gamma_t \left[\left(\lambda_2 + \frac{2\lambda_1}{\alpha_{q_t^*}} \right) \delta_{t,2} + \left(34 + \frac{8\lambda_1 \sigma(q_t^*)}{\sqrt{\alpha_{q_t^*}}} \right) \sqrt{\delta_{t,2}} \right], \end{aligned}$$

for any $p \in \mathcal{Q}_{\mathcal{X}}(1), q \in \mathcal{Q}_{\mathcal{Y}}(3)$ with probability at least $1 - \delta$ for sample size $M := \mathcal{O}(\epsilon^{-2} [\log(T\delta^{-1}) + d \log(1 + L_0 r_3(\epsilon, \lambda_1, \lambda_2) \epsilon^{-1})])$ in each iteration where $d := \max\{m, n\}$, α_p is the log-Sobolev coefficient of density p and $\sigma(p)$ denotes $(\mathbb{E}_p[\|\theta\|^2])^{1/2}$.

To obtain the complete convergence rate of the max-min gap, we need to specify the hyper-parameters τ_t, η_t, γ_t and sub-problem error bound $\delta_{t,1}, \delta_{t,2}$ for the minimax problem (2). By setting the input parameters as

$$\mu_t = \mu = \frac{1}{2} \left(2 + \lambda_2^2 - \lambda_2 \sqrt{4 + \lambda_2^2} \right), \quad \eta_t = \tau_t = \frac{\lambda_2 \mu}{1 - \mu}, \quad \gamma_t = \mu^{-t}, \quad (20)$$

and defining $c(p) = \min \left\{ \frac{1}{2}, \alpha_p, \frac{\alpha_p}{2(\sigma(p))^2} \right\}$, $c(q) = \min \left\{ \frac{1}{2}, \alpha_q, \frac{\alpha_q}{2(\sigma(q))^2} \right\}$ as the upper bound of the inner loop errors $\delta_{t,1}, \delta_{t,2}$, we are ready to provide the following outer loop error bound of our proposed algorithm with an inexact solver. It is worth noting that the parameter selection in (20) meets the pre-conditions of Theorem 8, in addition to satisfying $\mu \leq 1$ and $\mu^2 - (2 + \lambda_2^2)\mu + 1 = 0$. Moreover, noticing that MNE $w_* = (p_*, q_*)$ satisfies Eq. (16) and combining the optimality condition in Lemma 3, we can deduce that $p_* \in \mathcal{Q}_{\mathcal{X}}(1)$ and $q_* \in \mathcal{Q}_{\mathcal{Y}}(3)$. Therefore, we derive the following convergence result by combining Theorem 8.

Corollary 9. *Setting the parameters as Eq. (20) and assuming the sub-problem error bounds satisfy $\max\{\delta_{t,1}, \delta_{t,2}\} \leq \frac{1}{2} \min\{c(p_t^*), c(q_t^*)\} T^{-2J}$ for any fixed positive integer J and $\epsilon > 0$ at each step- t , the max-min gap between $\bar{w}_T := \frac{\sum_{t=1}^T \gamma_t w_t}{\sum_{t=1}^T \gamma_t}$ and $\bar{w}_T^* := \operatorname{argmax}_{w \in \mathcal{P}_2(\mathcal{X}) \times \mathcal{P}_2(\mathcal{Y})} Q(\bar{w}_T, w)$ satisfies*

$$0 \leq Q(\bar{w}_T, \bar{w}_T^*) \leq \frac{\lambda_2 \mu^T}{1 - \mu} [\text{KL}(\bar{p}_T^* \| p_0) + \text{KL}(\bar{q}_T^* \| q_0)] + 8\epsilon + (10\lambda_1 + \lambda_2 + 32)T^{-J}. \quad (21)$$

with probability at least $1 - \delta$ for sample size

$$M = \mathcal{O} \left(\epsilon^{-2} [\log(T\delta^{-1}) + d \log(1 + L_0 r_3(\epsilon, \lambda_1, \lambda_2) \epsilon^{-1})] \right),$$

in each iteration where $d = \max\{m, n\}$, $\bar{w}_T = \frac{\sum_{t=1}^T \gamma_t w_t}{\sum_{t=1}^T \gamma_t}$. Moreover, our high probability convergence also holds under the KL-divergence

$$\begin{aligned} \text{KL}(p_* \| p_T^*) + \frac{1}{2} \text{KL}(q_* \| q_T^*) &\leq \mu^T [\text{KL}(q_* \| q_0^*) + \text{KL}(p_* \| p_0^*)] \\ &\quad + \lambda_2^{-1} [8\epsilon + (10\lambda_1 + \lambda_2 + 32)T^{-J}], \end{aligned} \quad (22)$$

and the W_2 distance as follows

$$\begin{aligned} \frac{1}{2} W_2^2(q_T, q_*) + W_2^2(p_T, p_*) &\leq \frac{4}{\min\{\alpha_{p_T^*}, \alpha_{q_T^*}\}} \left\{ \mu^T [\text{KL}(q_* \| q_0^*) + \text{KL}(p_* \| p_0^*)] \right. \\ &\quad \left. + \lambda_2^{-1} [8\epsilon + (10\lambda_1 + \lambda_2 + 32)T^{-J}] \right\} + 3T^{-2J}, \end{aligned}$$

with respect to the MNE $w_* = (p_*, q_*)$ of Eq. (2).

6.2.2 INNER LOOP ERROR BOUND

In the preceding sub-section, we made the assumption that the KL divergence between (q_t, p_t) and (q_t^*, p_t^*) can be upper bounded by $(\delta_{t,1}, \delta_{t,2})$. In this sub-section, we provide evidence that the ULA algorithm effectively constrains the inner error bound $\text{KL}(p_t \| p_t^*), \text{KL}(q_t \| q_t^*)$

to a sufficiently small value (i.e. $\delta_{t,1}, \delta_{t,2}$) within polynomial time. We demonstrate that the summation of a quadratic function and a bounded function exhibits the log-Sobolev property. Similar to PDA (Nitanda et al., 2021), we utilize the Holley-Strook argument to estimate the log-Sobolev constant. Consequently, ULA converges efficiently to the target distribution, leveraging the log-Sobolev property. There are also other ways of estimating the log-Sobolev constants, such as Miclo’s trick. However, when leveraging the Lipschitz continuity, the log-Sobolev constant has exponential dependency on both dimension and Lipschitz constant. We focus on the the bounded condition for better computation complexity.

Lemma 10. (Holley and Stroock, 1987) Consider a probability density $p(\theta)$ on \mathbb{R}^{d_θ} satisfying the log-Sobolev inequality with a constant α . For a bounded function $H : \mathbb{R}^{d_\theta} \rightarrow \mathbb{R}$, we let a probability density $p_H(\theta)$ denote $p_H(\theta) := \frac{\exp(H(\theta))p(\theta)}{\mathbb{E}_p[\exp(H(\theta))]}$. Then, $p_H(\theta)$ satisfies the log-Sobolev inequality with $\alpha / \exp(4\|H\|_\infty)$.

It is clear that Eq. (14) and (15) can be rewritten as a summation of a bounded function and quadratic function. Lemma 10 shows that these two sub-problem can be log-Sobolev. For any $t \geq 0$, we can compute that the log-Sobolev constant for q_t^*, p_t^* are $\frac{\lambda_1}{\lambda_2} \exp\left(-\frac{12}{\lambda_2}\right)$, $\frac{\lambda_1}{\lambda_2} \exp\left(-\frac{4}{\lambda_2}\right)$, respectively, only depending on the choice of λ_1 and λ_2 (Refer to Appendix Lemma 28 for the constant estimation). The next theorem shows the convergence rate.

Theorem 11. (Vempala and Wibisono, 2019) Assume that $f(\theta)$ is smooth and has L_1 -Lipschitz gradient, and consider a probability density $p(\theta) \propto \exp(-f(\theta))$ satisfying the log-Sobolev inequality with constant α , ULA with step size $0 < \iota < \frac{\alpha}{4L_1^2}$ satisfies $\text{KL}(\rho_k \| p) \leq e^{-\alpha k} \text{KL}(\rho_0 \| p) + \frac{8\iota d_\theta L_1^2}{\alpha}$, where d_θ is the dimension.

6.2.3 GLOBAL CONVERGENCE

By leveraging Corollary 9 and Theorem 11, we are now ready to provide the global convergence of PAPAL. The ensuing theorem establishes the global error bound, achieved through the acquisition of a pair of $\frac{1}{2} \min\{c(p_t^*), c(q_t^*)\} T^{-2J}$ -approximate sub-problem solutions (p_t, q_t) under metrics $\text{KL}(\cdot \| p_t^*)$ and $\text{KL}(\cdot \| q_t^*)$ at each step- t , along the trajectory of ULA.

Theorem 12. Let ϵ be the desired accuracy and set the parameters $\gamma_t, \eta_t, \tau_t, \mu_t$ as in Eq. (20). Assuming Assumptions $[\mathbf{A}_1]$, $[\mathbf{A}_2]$ and $[\mathbf{A}_3]$, executing Langevin algorithm with step size $\iota = \mathcal{O}\left(\min\left\{1, \frac{\lambda_1^2}{\lambda_2^2}\right\} \frac{\lambda_1 \lambda_2 \epsilon^2}{d^2(L_1 + \lambda_1)^2 \exp(30/\lambda_2)}\right)$ for $T_t = \mathcal{O}\left(\iota^{-1} \lambda_1^{-1} [(\lambda_2 + 1) \exp(30/\lambda_2) \log(\epsilon^{-1})]\right)$ iterations on the inner loop yeilds an ϵ -mixed Nash equilibrium: $\max_{w \in \mathcal{P}_2(\mathcal{X}) \times \mathcal{P}_2(\mathcal{Y})} Q(\bar{w}, w) \leq \epsilon$ with probability at least $1 - \delta$ when running Algorithm 1 for iterations given by

$$T = \mathcal{O}\left(\max\left\{\log\left((1 + \lambda_2^{-1})\epsilon^{-1}\right) g^{-1}(\lambda_2), (32 + 10\lambda_1 + \lambda_2)\epsilon^{-1/J}\right\}\right),$$

and sample size $M = \mathcal{O}\left(\epsilon^{-2} [\log(T\delta^{-1}) + d \log(1 + L_0 r_3(\epsilon, \lambda_1, \lambda_2)\epsilon^{-1})]\right)$ on the outer loop, where $d = \max\{m, n\}$.

Then we obtain the computation time $T = \tilde{\mathcal{O}}(\max\{\log(\epsilon^{-1})\lambda_2^{-1}, \epsilon^{-1/J}\})$ if $\lambda_2 \rightarrow 0$, which indicates the limiting behavior of vanishing entropy regularizer. The dependence on λ_2^{-1} in

our convergence rate is a direct consequence of the classical LSI perturbation lemma (Holley and Stroock, 1987), which is likely unavoidable for Langevin-based methods in the most general setting (Menz and Schlichting, 2014). Moreover, according to Eq. (19), theorem 12 can also control $\text{KL}(\bar{p}||\bar{p}^*)$ and $\text{KL}(\bar{q}||\bar{q}^*)$ similarly. In addition, based on Corollary 9, we can also propose the following global convergence under metrics $\text{KL}(p_*||\cdot)$, $\text{KL}(q_*||\cdot)$ and $W_2(\cdot, p_*)$, $W_2(\cdot, q_*)$ with respect to MNE $w_* = (p_*, q_*)$ of Eq. (2).

Corollary 13. *Let ϵ be the desired accuracy. Under Assumptions $[\mathbf{A}_1]$, $[\mathbf{A}_2]$ and $[\mathbf{A}_3]$, if we execute ULA with a step size of $\iota = \mathcal{O}\left(\min\left\{\frac{\lambda_1}{\lambda_2}, \frac{\lambda_1^3}{\lambda_2^3}\right\} \frac{\lambda_2^4 \epsilon^2}{(\max\{m, n\})^2 (L_1 + \lambda_1)^2 \exp(30/\lambda_2)}\right)$ for $T_t = \mathcal{O}\left(\iota^{-1} \lambda_1^{-1} [(\lambda_2 + 1) \exp(30/\lambda_2) \log(\lambda_2^{-1} \epsilon^{-1})]\right)$ iterations on the inner loop, we can approximate an ϵ -accurate KL divergence: $\text{KL}(p_*||p_T^*) + \frac{1}{2}\text{KL}(q_*||q_T^*) \leq \epsilon$ with probability no less than $1 - \delta$ when running Algorithm 1 for iterations T and sample size M satisfy*

$$T = \mathcal{O}\left(\max\left\{\log\left((\lambda_2 \epsilon)^{-1}\right) g^{-1}(\lambda_2), (32 + 10\lambda_1 + \lambda_2)(\lambda_2 \epsilon)^{-1/J}\right\}\right),$$

$$M = \mathcal{O}\left(\lambda_2^{-2} \epsilon^{-2} \left[\log(T\delta^{-1}) + d \log(1 + L_0 r_3(\lambda_2 \epsilon, \lambda_1, \lambda_2))\right]\right),$$

on the outer loop with respect to the MNE $w_* = (p_*, q_*)$ of (2), where $d = \max\{m, n\}$. Furthermore, noting that p_T and q_T satisfying $\max\{\text{KL}(p_T||p_T^*), \text{KL}(q_T||q_T^*)\} \leq \mathcal{O}(\lambda_2^2 \epsilon^2)$, we can also approximate an ϵ -accurate W_2 distance: $W_2^2(p_T, p_*) + \frac{1}{2}W_2^2(q_T, q_*) \leq \epsilon$ with similar high probability when running Algorithm 1 for iterations T and sample size M satisfy

$$T = \mathcal{O}\left(\max\left\{\log\left((\tilde{g}(\lambda_1, \lambda_2)\epsilon)^{-1}\right) g^{-1}(\lambda_2), (32 + 10\lambda_1 + \lambda_2)(\tilde{g}(\lambda_1, \lambda_2)\epsilon)^{-1/J}\right\}\right),$$

$$M = \mathcal{O}\left(\tilde{g}^{-2}(\lambda_1, \lambda_2)\epsilon^{-2} \left[\log(T\delta^{-1}) + d \log(1 + L_0 r_3(\tilde{g}(\lambda_1, \lambda_2)\epsilon, \lambda_1, \lambda_2))\right]\right),$$

on the outer loop where $\tilde{g}(\lambda_1, \lambda_2) := \min\{1, \lambda_2, \lambda_1(\lambda_2)^{-1} \exp\{-12(\lambda_2)^{-1}\}\}$.

Inspired by MALA-based sampling algorithms (Chewi et al., 2021; Wu et al., 2022; Altschuler and Chewi, 2024), we also provide the global convergence of PAPAL using proximal sampler with MALA (refer to Algorithm 3) for inner loop.

Corollary 14. *Let ϵ be the desired accuracy and set the parameters $\gamma_t, \eta_t, \tau_t, \mu_t$ as in Eq. (20). Assuming Assumptions $[\mathbf{A}_1]$, $[\mathbf{A}_2]$ and $[\mathbf{A}_3]$, executing proximal sampler with MALA (Algorithm 3) for $T_t = \tilde{\mathcal{O}}\left((3L_1 + \lambda_1) \exp(12/\lambda_2) \lambda_1^{-1} d^{1/2} (\log(\epsilon^{-1}) + 1 + \lambda_2^{-1})^4\right)$ iterations on the inner loop yeilds an ϵ -mixed MNE: $\max_{w \in \mathcal{P}_2(\mathcal{X}) \times \mathcal{P}_2(\mathcal{Y})} Q(\bar{w}, w) \leq \epsilon$ with probability at least $1 - \delta$ when running Algorithm 1 for iterations given by*

$$T = \mathcal{O}\left(\max\left\{\log\left((1 + \lambda_2^{-1})\epsilon^{-1}\right) g^{-1}(\lambda_2), (32 + 10\lambda_1 + \lambda_2)\epsilon^{-1/J}\right\}\right),$$

and sample size $M = \mathcal{O}\left(\epsilon^{-2} \left[\log(T\delta^{-1}) + d \log(1 + L_0 r_3(\epsilon, \lambda_1, \lambda_2)\epsilon^{-1})\right]\right)$ on the outer loop, where $d = \max\{m, n\}$.

6.3 Convergence of Stochastic PAPAL

In this subsection, we consider the stochastic version of Eq.(2) as follows

$$\min_{p \in \mathcal{P}_2(\mathcal{X})} \max_{q \in \mathcal{P}_2(\mathcal{Y})} \mathbb{E}_{p(x)} \mathbb{E}_{q(y)} [f(x, y) := \mathbb{E}[G(x, y, \xi)]] + \mathcal{R}(p) - \mathcal{R}(q), \quad (23)$$

where $\mathcal{X} = \mathbb{R}^m$, $\mathcal{Y} = \mathbb{R}^n$, and the stochastic component $G(x, y, \xi)$ is indexed by some random vector ξ whose probability distribution ρ is supported on $\Lambda \subset \mathbb{R}^d$ and $G : \mathcal{X} \times \Lambda \rightarrow \mathbb{R}$. We assume that the expectation

$$\mathbb{E}[G(x, y, \xi)] = \int_{\Lambda} G(x, y, \xi) d\rho(\xi), \quad (24)$$

is well defined and finite valued for every $(x, y) \in \mathcal{X} \times \mathcal{Y}$. To prevent ambiguity, we will continue to use $\mathcal{L}(p, q)$ to denote $\mathbb{E}_{p(x)} \mathbb{E}_{q(y)} [f(x, y)] + \mathcal{R}(p) - \mathcal{R}(q)$.

For stochastic version, we have additional assumptions,

[A₄] There exists L_0 such that $G(\cdot, \cdot, \xi)$ is L_0 -Lipschitz for any $\xi \in \Lambda$.

[A₅] For any $\xi \in \Lambda$, $G(\cdot, \cdot, \xi)$ has L_1 -Lipschitz gradient.

[A₆] For any $x, y \in \mathcal{X} \times \mathcal{Y}$ and $\xi \in \Lambda$, $|G(x, y, \xi)| \leq 1$.

To accommodate the problem setting Eq. (23), we need to make some straightforward modifications to the deterministic PAPAL algorithm: we replace function $\hat{\phi}_t$ with $(\hat{\phi}_t)_{\xi_N} := -(1 + \mu_t) \mathbb{E}_{\hat{p}_{t-1}} \left[\frac{1}{N} \sum_{i=1}^N G(\cdot, y, \xi_{t,i}) \right] + \mu_t \mathbb{E}_{\hat{p}_{t-2}} \left[\frac{1}{N} \sum_{i=1}^N G(\cdot, y, \xi_{t,i}) \right]$, and $\hat{\psi}_t$ with $(\hat{\psi}_t)_{\xi_N} := \mathbb{E}_{\hat{q}_t} \left[\frac{1}{N} \sum_{i=1}^N G(\cdot, y, \xi_{t,i}) \right]$ in each iteration t of the Algorithm 1, where $(\xi_{t,1}, \dots, \xi_{t,N})$ denote N i.i.d random variables, depending on t -th update with $\xi_{t,i} = \xi$ for any $i \in [1 : N]$.

6.3.1 OUTER LOOP ERROR BOUND

In order to simplify the notation of our formula, we also assume that the KL divergence between (q_t, p_t) and (q_t^*, p_t^*) can be bounded by small errors $(\delta_{t,1}, \delta_{t,2})$ at each step- t . As an extension of Corollary 9 in the stochastic version, the following lemma provides the high-probability global convergence of the stochastic PAPAL algorithm under the condition of sufficiently small inner loop error.

Lemma 15. *Setting the parameters $\gamma_t, \eta_t, \tau_t, \mu_t$ as in Eq. (20) and assuming the sub-problem error bounds satisfy $\max\{\delta_{t,1}, \delta_{t,2}\} \leq \frac{1}{2} \min\{c(p_t^*), c(q_t^*)\} T^{-2J}$ for any fixed positive integer J and $\epsilon > 0$ at each step- t . If Assumptions [A₄], [A₅] and [A₆] hold, then the max-min gap between \bar{w}_T and $\bar{w}_T^* := \operatorname{argmax}_{w \in \mathcal{P}_2(\mathcal{X}) \times \mathcal{P}_2(\mathcal{Y})} Q(\bar{w}_T, w)$ satisfies*

$$0 \leq Q(\bar{w}_T, \bar{w}_T^*) \leq \frac{\lambda_2 \mu^T}{1 - \mu} [\text{KL}(\bar{p}_T^* \| p_0) + \text{KL}(\bar{q}_T^* \| q_0)] + 12\epsilon + (10\lambda_1 + \lambda_2 + 32)T^{-J}. \quad (25)$$

with probability at least $1 - \delta$ for sample size

$$M = N = \mathcal{O} \left(\epsilon^{-2} \left[\log(2T\delta^{-1}) + d \log(1 + L_0 r_3(\epsilon, \lambda_1, \lambda_2) \epsilon^{-1}) \right] \right),$$

in each iteration where $d = \max\{m, n\}$, $\bar{w}_T = \frac{\sum_{t=1}^T \gamma_t w_t}{\sum_{t=1}^T \gamma_t}$. In addition, our high probability convergence also holds under the KL-divergence

$$\begin{aligned} \text{KL}(p_* \| p_T^*) + \frac{1}{2} \text{KL}(q_* \| q_T^*) &\leq \mu^T [\text{KL}(q_* \| q_0^*) + \text{KL}(p_* \| p_0^*)] \\ &\quad + \lambda_2^{-1} [12\epsilon + (10\lambda_1 + \lambda_2 + 32)T^{-J}], \end{aligned} \quad (26)$$

and the W_2 distance as follows

$$\begin{aligned} \frac{1}{2} W_2^2(q_T, q_*) + W_2^2(p_T, p_*) &\leq \frac{4}{\min\{\alpha_{p_T^*}, \alpha_{q_T^*}\}} \left\{ \mu^T [\text{KL}(q_* \| q_0^*) + \text{KL}(p_* \| p_0^*)] \right. \\ &\quad \left. + \lambda_2^{-1} [12\epsilon + (10\lambda_1 + \lambda_2 + 32)T^{-J}] \right\} + 3T^{-2J}, \end{aligned}$$

with respect to the MNE $w_* = (p_*, q_*)$ of Eq. (23).

For the inner loop, the stochastic version is equivalent to the deterministic one (refer 6.2.2). Thus, we omit the discussion of the inner loop error bound for Stochastic PAPAL.

6.3.2 GLOBAL CONVERGENCE

Theorem 16. *Let ϵ be the desired accuracy. Under Assumptions $[\mathbf{A}_4]$, $[\mathbf{A}_5]$ and $[\mathbf{A}_6]$, if we run Langevin algorithm with step size $\iota = \mathcal{O}\left(\min\left\{1, \frac{\lambda_1^2}{\lambda_2^2}\right\} \frac{\lambda_1 \lambda_2 \epsilon^2}{d^2(L_1 + \lambda_1)^2 \exp(30/\lambda_2)}\right)$ for $T_t = \mathcal{O}\left(\iota^{-1} \lambda_1^{-1} [(\lambda_2 + 1) \exp(30/\lambda_2) \log(\epsilon^{-1})]\right)$ iterations on the inner loop, then we can achieve an ϵ -mixed Nash equilibrium: $\max_{w \in \mathcal{P}_2(\mathcal{X}) \times \mathcal{P}_2(\mathcal{Y})} Q(\bar{w}, w) \leq \epsilon$ with probability at least $1 - \delta$ when running Algorithm 1 for iterations*

$$T = \mathcal{O}\left(\max\left\{\log\left((1 + \lambda_2^{-1})\epsilon^{-1}\right) g^{-1}(\lambda_2), (32 + 10\lambda_1 + \lambda_2)\epsilon^{-1/J}\right\}\right),$$

and sample size

$$M = N = \mathcal{O}\left(\epsilon^{-2} [\log(2T\delta^{-1}) + d \log(1 + L_0 r_3(\epsilon, \lambda_1, \lambda_2)\epsilon^{-1})]\right),$$

on the outer loop, where $d = \max\{m, n\}$.

Similar to the conclusion of Corollary 13 in the deterministic setting, the stochastic PAPAL algorithm also exhibits global convergence properties with respect to the KL divergence and W_2 distance of MNE $w_* = (p_*, q_*)$.

Corollary 17. *Let ϵ be the desired accuracy. Under Assumptions $[\mathbf{A}_4]$, $[\mathbf{A}_5]$ and $[\mathbf{A}_6]$, if we execute ULA with a step size of $\iota = \mathcal{O}\left(\min\left\{\frac{\lambda_1}{\lambda_2}, \frac{\lambda_1^3}{\lambda_2^3}\right\} \frac{\lambda_2^4 \epsilon^2}{(\max\{m, n\})^2 (L_1 + \lambda_1)^2 \exp(30/\lambda_2)}\right)$ for $T_t = \mathcal{O}\left(\iota^{-1} \lambda_1^{-1} [(\lambda_2 + 1) \exp(30/\lambda_2) \log(\lambda_2^{-1} \epsilon^{-1})]\right)$ iterations on the inner loop, we can approximate an ϵ -accurate KL divergence: $\text{KL}(p_* \| p_T^*) + \frac{1}{2} \text{KL}(q_* \| q_T^*) \leq \epsilon$ with probability no less than $1 - \delta$ when running Algorithm 1 for iterations T and sample size M, N satisfy*

$$T = \mathcal{O}\left(\max\left\{\log\left((\lambda_2 \epsilon)^{-1}\right) g^{-1}(\lambda_2), (32 + 10\lambda_1 + \lambda_2)(\lambda_2 \epsilon)^{-1/J}\right\}\right),$$

$$M = N = \mathcal{O}\left(\lambda_2^{-2} \epsilon^{-2} [\log(T\delta^{-1}) + d \log(1 + L_0 r_3(\lambda_2 \epsilon, \lambda_1, \lambda_2))]\right),$$

on the outer loop with respect to the MNE $w_* = (p_*, q_*)$ of (2), where $d = \max\{m, n\}$. Furthermore, noting that p_T and q_T satisfying $\max\{\text{KL}(p_T \| p_T^*), \text{KL}(q_T \| q_T^*)\} \leq \mathcal{O}(\lambda_2^2 \epsilon^2)$, we can also approximate an ϵ -accurate W_2 distance: $W_2^2(p_T, p_*) + \frac{1}{2}W_2^2(q_T, q_*) \leq \epsilon$ with similar high probability when running Algorithm 1 for iterations T and sample size M, N satisfy

$$T = \mathcal{O} \left(\max \left\{ \log \left((\tilde{g}(\lambda_1, \lambda_2) \epsilon)^{-1} \right) g^{-1}(\lambda_2), (32 + 10\lambda_1 + \lambda_2) (\tilde{g}(\lambda_1, \lambda_2) \epsilon)^{-1/J} \right\} \right),$$

$$M = N = \mathcal{O} \left(\tilde{g}^{-2}(\lambda_1, \lambda_2) \epsilon^{-2} \left[\log(T\delta^{-1}) + d \log(1 + L_0 r_3(\tilde{g}(\lambda_1, \lambda_2) \epsilon, \lambda_1, \lambda_2)) \right] \right),$$

on the outer loop where $\tilde{g}(\lambda_1, \lambda_2) := \min \{1, \lambda_2, \lambda_1(\lambda_2)^{-1} \exp\{-12(\lambda_2)^{-1}\}\}$.

Although stochastic PAPAL needs large batch size $\Theta(\epsilon^{-2})$ for the outer loop, the inner loop can also be implemented with stochastic algorithms, such as some SGLD-based algorithms (Freund et al., 2022; Das et al., 2023; Huang et al., 2024). According to the result of Das et al. (2023), we can extend our algorithm with batch size 1. We begin the discussion with the following assumption:

Assumption 6.1 (Sub-Gaussian Stochastic Gradient Growth). *For $L_0 > 0$, the function $G(x, y, \xi)$ satisfies the following norm-subgaussianity condition for the stochastic gradient*

$$\mathbb{P} \{ \|\nabla G(x, y, \xi) - \nabla f(x, y)\| \geq t \mid x, y \} \leq 2 \exp \left\{ -t^2 / (8L_0^2) \right\}, \forall (x, y) \in \mathcal{X} \times \mathcal{Y}.$$

Therefore, applying Das et al. (2023, Theorem 5), we have following estimation of sub-problem

$$\begin{aligned} \max \{ \text{KL}(p_t \| p_t^*), \text{KL}(q_t \| q_t^*) \} &\leq \exp\{-\lambda \iota T_t / 2\} \max \{ \text{KL}(p_0 \| p_t^*), \text{KL}(q_0 \| q_t^*) \} \\ &+ \mathcal{O} \left(\frac{(3L_1 + \lambda_1)^2 \iota d}{\lambda \lambda_2^2} + \frac{M \iota (3L_1 + \lambda_1) L_0^2}{\lambda \lambda_2^3 N} + \frac{M \iota L_0^4}{\lambda \lambda_2^4 N^2} \right), \end{aligned}$$

using SGLD with arbitrary batch size $N \in \mathbb{N}_+$ in each iteration t of PAPAL algorithm, where M is the number of particles, T_t is the number of iterations on the inner loop, ι is step size, $\lambda = \lambda_1 \exp\{-12\lambda_2^{-1}\} \lambda_2^{-1}$ and $d = \max\{m, n\}$.

Corollary 18. *Let ϵ be the desired accuracy set the parameters $\gamma_t, \eta_t, \tau_t, \mu_t$ as default. Under Assumptions $[\mathbf{A}_4]$, $[\mathbf{A}_5]$, $[\mathbf{A}_6]$ and subgaussian stochastic gradient growth condition 6.1, if we execute SGLD with batch size $N \in \mathbb{N}_+$ and a step size of*

$$\iota = \mathcal{O} \left(\frac{\lambda_1}{\lambda_2} \min \{ \lambda_1^2, \lambda_2^2 \} \min \left\{ \frac{\lambda_2^2 N^2}{L_0^4}, \frac{1}{(3L_1 + \lambda_1)^2} \right\} \frac{\exp(-30/\lambda_2)}{\max\{M, d\}} \right)$$

for $T_t = \mathcal{O} \left(\iota^{-1} \lambda_1^{-1} [(\lambda_2 + 1) \exp(30/\lambda_2) \log(\epsilon^{-1})] \right)$ iterations on the inner loop yeilds an ϵ -mixed Nash equilibrium with probability at least $1 - \delta$ when running PAPAL for iteration given by

$$T = \mathcal{O} \left(\max \left\{ \log \left((1 + \lambda_2^{-1}) \epsilon^{-1} \right) g^{-1}(\lambda_2), (32 + 10\lambda_1 + \lambda_2) \epsilon^{-1/J} \right\} \right),$$

and sample size $M = \mathcal{O} \left(\epsilon^{-2} \left[\log(T\delta^{-1}) + d \log(1 + L_0 r_3(\epsilon, \lambda_1, \lambda_2) \epsilon^{-1}) \right] \right)$ on the outer loop, where $d = \max\{m, n\}$.

7. Conclusion

We propose a particle-based algorithm – PAPAL to solve min-max problems in continuous probability space. This algorithm explores the global convergence property for finding general mixed Nash equilibria in infinite-dimensional two-player zero-sum games. We provide a quantitative computation and sample complexity analysis of PAPAL with a suitable entropy regularization, which first shows the feasibility of an implementable particle-based algorithm in this task.

References

- Leonard Adolphs, Hadi Daneshmand, Aurelien Lucchi, and Thomas Hofmann. Local saddle point optimization: A curvature exploitation approach. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 486–495. PMLR, 2019.
- Jason M Altschuler and Sinho Chewi. Faster high-accuracy log-concave sampling via algorithmic warm starts. *Journal of the ACM*, 71(3):1–55, 2024.
- Sanjeev Arora, Rong Ge, Yingyu Liang, Tengyu Ma, and Yi Zhang. Generalization and equilibrium in generative adversarial nets (gans). In *International Conference on Machine Learning*, pages 224–232. PMLR, 2017.
- Francis Bach. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017.
- Dominique Bakry and Michel Émery. Diffusions hypercontractives. In *Séminaire de Probabilités XIX 1983/84: Proceedings*, pages 177–206. Springer, 2006.
- Maximilian Balandat, Walid Krichene, Claire Tomlin, and Alexandre Bayen. Minimizing regret on reflexive banach spaces and nash equilibria in continuous zero-sum games. *Advances in Neural Information Processing Systems*, 29, 2016.
- Tamer Başar and Geert Jan Olsder. *Dynamic noncooperative game theory*. SIAM, 1998.
- Yoshua Bengio, Nicolas Roux, Pascal Vincent, Olivier Delalleau, and Patrice Marcotte. Convex neural networks. *Advances in neural information processing systems*, 18, 2005.
- Sergey G Bobkov, Ivan Gentil, and Michel Ledoux. Hypercontractivity of hamilton–jacobi equations. *Journal de Mathématiques Pures et Appliquées*, 80(7):669–696, 2001.
- Kristian Bredies and Hongpeng Sun. Preconditioned douglas–rachford splitting methods for convex-concave saddle-point problems. *SIAM Journal on Numerical Analysis*, 53(1):421–444, 2015.
- Lucian Busoniu, Robert Babuska, and Bart De Schutter. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2):156–172, 2008.
- George HG Chen and R Tyrrell Rockafellar. Convergence rates in forward–backward splitting. *SIAM Journal on Optimization*, 7(2):421–444, 1997.

- Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Computing nash equilibria: Approximation and smoothed complexity. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 603–612. IEEE, 2006.
- Sinho Chewi, Chen Lu, Kwangjun Ahn, Xiang Cheng, Thibaut Le Gouic, and Philippe Rigollet. Optimal dimension dependence of the metropolis-adjusted langevin algorithm. In *Conference on Learning Theory*, pages 1260–1300. PMLR, 2021.
- Lénaïc Chizat. Mean-field langevin dynamics: Exponential convergence and annealing. *arXiv preprint arXiv:2202.01009*, 2022.
- Aniket Das, Dheeraj M Nagaraj, and Anant Raj. Utilising the clt structure in stochastic gradient based sampling: Improved analysis and faster algorithms. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 4072–4129. PMLR, 2023.
- Partha Dasgupta and Eric Maskin. The existence of equilibrium in discontinuous economic games, i: Theory. *The Review of economic studies*, 53(1):1–26, 1986.
- Constantinos Daskalakis. On the complexity of approximating a nash equilibrium. *ACM Transactions on Algorithms (TALG)*, 9(3):1–35, 2013.
- Constantinos Daskalakis and Ioannis Panageas. The limit points of (optimistic) gradient descent in min-max optimization. *Advances in neural information processing systems*, 31, 2018.
- Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with optimism. *arXiv preprint arXiv:1711.00141*, 2017.
- Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The complexity of constrained min-max optimization. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1466–1478, 2021.
- Jelena Diakonikolas, Constantinos Daskalakis, and Michael I Jordan. Efficient methods for structured nonconvex-nonconcave min-max optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 2746–2754. PMLR, 2021.
- Roland L’vovich Dobrushin. Vlasov equations. *Funktsional’nyi Analiz i ego Prilozheniya*, 13(2):48–58, 1979.
- Carles Domingo-Enrich, Samy Jelassi, Arthur Mensch, Grant Rotskoff, and Joan Bruna. A mean-field analysis of two-player zero-sum games. *Advances in neural information processing systems*, 33:20215–20226, 2020.
- David L Donoho et al. High-dimensional data analysis: The curses and blessings of dimensionality. *AMS math challenges lecture*, 1(2000):32, 2000.
- Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid monte carlo. *Physics letters B*, 195(2):216–222, 1987.

- Yu G Evtushenko. Some local properties of minimax problems. *USSR Computational Mathematics and Mathematical Physics*, 14(3):129–138, 1974.
- Cong Fang, Jason Lee, Pengkun Yang, and Tong Zhang. Modeling from features: a mean-field framework for over-parameterized deep neural networks. In *Conference on learning theory*, pages 1887–1936. PMLR, 2021.
- Cong Fang, Yihong Gu, Weizhong Zhang, and Tong Zhang. Convex formulation of overparameterized deep neural networks. *IEEE Transactions on Information Theory*, 2022.
- Tanner Fiez, Benjamin Chasnov, and Lillian J Ratliff. Convergence of learning dynamics in stackelberg games. *arXiv preprint arXiv:1906.01217*, 2019.
- Yoav Freund, Yi-An Ma, and Tong Zhang. When is the convergence time of langevin algorithms dimension independent? a composite optimization viewpoint. *Journal of Machine Learning Research*, 23(214):1–32, 2022.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.
- Irving L Glicksberg. A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Ulf Grenander and Michael I Miller. Representations of knowledge in complex systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 56(4):549–581, 1994.
- Paulina Grnarova, Kfir Y Levy, Aurelien Lucchi, Thomas Hofmann, and Andreas Krause. An online learning approach to generative adversarial networks. *arXiv preprint arXiv:1706.03269*, 2017.
- Leonard Gross. Logarithmic sobolev inequalities. *American Journal of Mathematics*, 97(4):1061–1083, 1975.
- Bingsheng He, Feng Ma, and Xiaoming Yuan. An algorithmic framework of generalized primal-dual hybrid gradient methods for saddle point problems. *Journal of Mathematical Imaging and Vision*, 58(2):279–293, 2017.
- Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
- Richard Holley and Daniel W Stroock. Logarithmic sobolev inequalities and stochastic ising models. *Journal of Statistical Physics*, 46, 1987.

- Ya-Ping Hsieh, Chen Liu, and Volkan Cevher. Finding mixed nash equilibria of generative adversarial networks. In *International Conference on Machine Learning*, pages 2810–2819. PMLR, 2019.
- Kaitong Hu, Zhenjie Ren, David Šiška, and Lukasz Szpruch. Mean-field langevin dynamics and energy landscape of neural networks. In *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, volume 57, pages 2043–2065. Institut Henri Poincaré, 2021.
- Xunpeng Huang, Difan Zou, Hanze Dong, Yian Ma, and Tong Zhang. Faster sampling via stochastic gradient proximal sampler. In *Forty-first International Conference on Machine Learning*, 2024.
- Chi Jin, Praneeth Netrapalli, and Michael Jordan. What is local optimality in nonconvex-nonconcave minimax optimization? In *International conference on machine learning*, pages 4880–4889. PMLR, 2020.
- Galina M Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.
- Guanghui Lan. *First-order and stochastic optimization methods for machine learning*. Springer, 2020.
- Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- Tianyi Lin, Chi Jin, and Michael Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In *International Conference on Machine Learning*, pages 6083–6093. PMLR, 2020.
- Songtao Lu, Ioannis Tsaknakis, Mingyi Hong, and Yongxin Chen. Hybrid block successive approximation for one-sided non-convex min-max problems: algorithms and applications. *IEEE Transactions on Signal Processing*, 68:3676–3691, 2020.
- Yulong Lu. Two-scale gradient descent ascent dynamics finds mixed nash equilibria of continuous games: A mean-field perspective. *CoRR*, abs/2212.08791, 2022.
- D Russell Luke and Ron Shefi. A globally linearly convergent method for pointwise quadratically supportable convex–concave saddle point problems. *Journal of Mathematical Analysis and Applications*, 457(2):1568–1590, 2018.
- Chao Ma and Lexing Ying. Provably convergent quasistatic dynamics for mean-field two-player zero-sum games. In *The Tenth International Conference on Learning Representations, ICLR*. OpenReview.net, 2022.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- Oren Mangoubi and Nisheeth K Vishnoi. Greedy adversarial equilibrium: an efficient alternative to nonconvex-nonconcave min-max optimization. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 896–909, 2021.

- Eric Mazumdar, Lillian J Ratliff, and S Shankar Sastry. On gradient-based learning in continuous games. *SIAM Journal on Mathematics of Data Science*, 2(1):103–131, 2020.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In *Conference on Learning Theory*, pages 2388–2464. PMLR, 2019.
- Georg Menz and André Schlichting. Poincaré and logarithmic Sobolev inequalities by decomposition of the energy landscape. *The Annals of Probability*, 42(5):1809 – 1884, 2014. doi: 10.1214/14-AOP908. URL <https://doi.org/10.1214/14-AOP908>.
- Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International conference on machine learning*, pages 3481–3490. PMLR, 2018.
- Hongseok Namkoong and John C Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. *Advances in neural information processing systems*, 29, 2016.
- Radford M Neal et al. Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2, 2011.
- Angelia Nedić and Asuman Ozdaglar. Subgradient methods for saddle-point problems. *Journal of optimization theory and applications*, 142(1):205–228, 2009.
- Arkadi Nemirovski. Prox-method with rate of convergence $o(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- Arkadi Nemirovski and D Yudin. On cezari’s convergence of the steepest descent method for approximating saddle point of convex-concave functions. In *Soviet Mathematics. Doklady*, volume 19, pages 258–269, 1978.
- Phan-Minh Nguyen and Huy Tuan Pham. A rigorous framework for the mean field limit of multilayer neural networks. *arXiv preprint arXiv:2001.11443*, 2020.
- Atsushi Nitanda, Denny Wu, and Taiji Suzuki. Particle dual averaging: Optimization of mean field neural network with global convergence rate analysis. *Advances in Neural Information Processing Systems*, 34:19608–19621, 2021.
- Shayegan Omidshafiei, Jason Pazis, Christopher Amato, Jonathan P How, and John Vian. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In *International Conference on Machine Learning*, pages 2681–2690. PMLR, 2017.

- Felix Otto and Cédric Villani. Generalization of an inequality by talagrand and links with the logarithmic sobolev inequality. *Journal of Functional Analysis*, 173(2):361–400, 2000.
- Hassan Rafique, Mingrui Liu, Qihang Lin, and Tianbao Yang. Non-convex min-max optimization: Provable algorithms and applications in machine learning. 2019.
- Peter J Rossky, Jimmie D Doll, and Harold L Friedman. Brownian dynamics as smart monte carlo simulation. *The Journal of Chemical Physics*, 69(10):4628–4633, 1978.
- Tim Roughgarden. Algorithmic game theory. *Communications of the ACM*, 53(7):78–86, 2010.
- Aviad Rubinfeld. Settling the complexity of computing approximate two-player nash equilibria. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 258–265. IEEE, 2016.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Aman Sinha, Hongseok Namkoong, Riccardo Volpi, and John Duchi. Certifying some distributional robustness with principled adversarial training. *arXiv preprint arXiv:1710.10571*, 2017.
- Michel Talagrand. Concentration of measure and isoperimetric inequalities in product spaces. *Publications Mathématiques de l’Institut des Hautes Etudes Scientifiques*, 81(1):73–205, 1995.
- Vladimir Vapnik. Principles of risk minimization for learning theory. *Advances in neural information processing systems*, 4, 1991.
- Santosh Vempala and Andre Wibisono. Rapid convergence of the unadjusted langevin algorithm: Isoperimetry suffices. *Advances in neural information processing systems*, 32, 2019.
- John Von Neumann. On the theory of parlor games. *Mathematische Annalen*, 100(1):295–320, 1928.
- John Von Neumann and Oskar Morgenstern. Theory of games and economic behavior, 2nd rev. 1947.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.
- Guillaume Wang and Lénaïc Chizat. An exponentially converging particle method for the mixed nash equilibrium of continuous games. *CoRR*, abs/2211.01280, 2022.

Keru Wu, Scott Schmidler, and Yuansi Chen. Minimax mixing time of the metropolis-adjusted langevin algorithm for log-concave sampling. *Journal of Machine Learning Research*, 23(270):1–63, 2022.

Greg Yang and Edward J Hu. Feature learning in infinite-width neural networks. *arXiv preprint arXiv:2011.14522*, 2020.

Greg Yang, Michael Santacroce, and Edward J Hu. Efficient computation of deep nonlinear infinite-width neural networks that learn features. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=tUMrOIox8XW>.

Appendix A. Auxiliary lemmas

Proposition 19. (Bakry and Émery, 2006) Consider a probability density $p(\theta) \propto \exp(-f(\theta))$, where $f: \mathbb{R}^{d_\theta} \rightarrow \mathbb{R}$ is a smooth function. If there exists $\alpha > 0$ such that $\nabla^2 f \succeq \alpha I_{d_\theta}$, then $p(\theta)$ satisfies log-Sobolev inequality with α .

Definition 20 (Talagrand’s inequality). We say that a probability density $p(\theta)$ satisfies Talagrand’s inequality (Talagrand, 1995) with a constant $\alpha > 0$ if for any probability density $p'(\theta)$, $\frac{\alpha}{2} W_2^2(p', p) \leq \text{KL}(p' \| p)$, where $W_2(p', p)$ denotes the 2-Wasserstein distance between $p'(\theta)$ and $p(\theta)$.

Lemma 21. (Otto and Villani, 2000; Bobkov et al., 2001) If a probability density $p(\theta)$ satisfies the log-Sobolev inequality with a constant $\alpha > 0$, then $p(\theta)$ satisfies Talagrand’s inequality with the same α .

Theorem 22. (Vempala and Wibisono, 2019) Assume $f(\theta)$ is smooth and has L_1 -Lipschitz gradient, and consider a probability density $p(\theta) \propto \exp(-f(\theta))$ satisfying the log-Sobolev inequality with constant α , ULA with step size $0 < \iota < \frac{\alpha}{4L_1^2}$ satisfies

$$\text{KL}(\rho_k \| p) \leq e^{-\alpha k} \text{KL}(\rho_0 \| p) + \frac{8\iota d_\theta L_1^2}{\alpha}. \quad (27)$$

Lemma 23 (Tail bound for Chi-squared variable Laurent and Massart (2000)). Let $\theta \sim \mathcal{N}(0, \sigma^2 I_{d_\theta})$ be a Gaussian random variable on \mathbb{R}^{d_θ} . Then, we obtain

$$\mathbb{P}[\|\theta\|^2 \geq 2c] \leq \exp\left(-\frac{c}{10\sigma^2}\right), \quad \forall c \geq d_\theta \sigma^2. \quad (28)$$

Lemma 24 (Hoeffding’s inequality (Hoeffding, 1963)). Let X, X_1, \dots, X_M be i.i.d random variables taking values in $[-a, a]$ for $a > 0$. Then, for any $\rho > 0$, we get

$$\mathbb{P}\left[\left|\frac{1}{M} \sum_{r=1}^M X_r - \mathbb{E}[X]\right| > \rho\right] \leq 2 \exp\left(-\frac{\rho^2 M}{2a^2}\right). \quad (29)$$

Appendix B. Existence of Mixed Nash Equilibrium

Proof [Proof of Theorem 6] For the purpose of ensuring a seamless proof, it is necessary to introduce the following lemma.

Lemma 25. *The following set*

$$\mathcal{P}(g, \mathcal{L}, \mathbb{R}^d) := \left\{ p : \mathbb{R}^d \rightarrow \mathbb{R}^+ \mid p \leq g \text{ a.e. on } \mathbb{R}^d, p \in \mathcal{B}(\mathcal{L}), \int_{\mathbb{R}^d} p(x) dx = 1 \right\}, \quad (30)$$

is a convex compact set where

$$\mathcal{B}(\mathcal{L}) := \left\{ f : \mathbb{R}^d \rightarrow \mathbb{R} \mid \max_{x, y \in B_k(\mathbf{0})} \frac{|f(x) - f(y)|}{\|x - y\|} \leq L_k, \forall k \in \mathbb{N}^+ \right\},$$

$g \in L^1$ and $\mathcal{L} := \{L_k\}_{k=1}^\infty$.

Proof It is easy to notice that $\mathcal{P}(g, \mathcal{L}, \mathbb{R}^d)$ is convex. Considering the infinite sets $\{p_n\}_{n=1}^\infty \subset \mathcal{P}(g, \mathcal{L}, \mathbb{R}^d)$ and $\{B_n(\mathbf{0}) \subset \mathbb{R}^d\}_{n=1}^\infty$, we proceed to construct a subsequence of $\{p_n\}_{n=1}^\infty$ in the following manner:

1. Applying the Arzelà–Ascoli theorem, we ascertain the existence of a subsequence $\{p_{1n}\}_{n=1}^\infty$ of $\{p_n\}_{n=1}^\infty$, which converges uniformly to $p^*|_{B_1(\mathbf{0})}$ on $B_1(\mathbf{0})$. Moreover, $p^*|_{B_1(\mathbf{0})}$ satisfies the Lipschitz condition:

$$\max_{x, y \in B_1(\mathbf{0})} \frac{|p^*|_{B_1(\mathbf{0})}(x) - p^*|_{B_1(\mathbf{0})}(y)|}{\|x - y\|} \leq L_1.$$

2. For $k \in [1 : \infty)$, we employ the Arzelà–Ascoli theorem once again to obtain a subsequence $\{p_{(k+1)n}\}_{n=1}^\infty$ of $\{p_{kn}\}_{n=k}^\infty$, which converges uniformly to $p^*|_{B_{k+1}(\mathbf{0})}$ on $B_{k+1}(\mathbf{0})$. Furthermore, $p^*|_{B_{k+1}(\mathbf{0})}$ satisfies the Lipschitz condition:

$$\max_{x, y \in B_{k+1}(\mathbf{0})} \frac{|p^*|_{B_{k+1}(\mathbf{0})}(x) - p^*|_{B_{k+1}(\mathbf{0})}(y)|}{\|x - y\|} \leq L_{k+1}.$$

As a result, we can establish the existence of a subsequence $\{p_{kk}\}_{k=1}^\infty$ of $\{p_n\}_{n=1}^\infty$, which converges to p^* a.e. on \mathbb{R}^d . Noting $g(x)$ belongs to L^1 , then it's intuitive that $p^* \in L^1$ and $\|p_{kk} - p^*\|_{L^1} \rightarrow 0$ by using dominate convergence theorem, which indicates that $p^* \in \mathcal{P}(g, \mathcal{L}, \mathbb{R}^d)$. \blacksquare

Note that the following two expressions are equivalent

- 1.

$$\hat{p} = \operatorname{argmin}_{p \in \mathcal{P}_2(\mathbb{R}^m)} \mathbb{E}_{p(x)} \mathbb{E}_{\bar{q}(y)} [f(x, y)] + \mathcal{R}(p) - \mathcal{R}(\bar{q}), \quad (31)$$

$$\hat{q} = \operatorname{argmin}_{q \in \mathcal{P}_2(\mathbb{R}^n)} \mathbb{E}_{\bar{p}(x)} \mathbb{E}_{q(y)} [f(x, y)] + \mathcal{R}(\bar{p}) - \mathcal{R}(q). \quad (32)$$

2.

$$\hat{p} \propto \exp \left\{ -(\lambda_2)^{-1} \left[\mathbb{E}_{\bar{q}(y)} [f(x, y)] + \lambda_1 \|x\|^2 \right] \right\}, \quad (33)$$

$$\hat{q} \propto \exp \left\{ -(\lambda_2)^{-1} \left[-\mathbb{E}_{\bar{p}(x)} [f(x, y)] + \lambda_1 \|y\|^2 \right] \right\}, \quad (34)$$

for any $(\bar{p}, \bar{q}) \in \mathcal{P}_2(\mathbb{R}^m) \times \mathcal{P}_2(\mathbb{R}^n)$. Therefore, we denote operator $\mathbf{F}(p, q) : \mathcal{P}_2(\mathbb{R}^m) \times \mathcal{P}_2(\mathbb{R}^n) \rightarrow \mathcal{P}_2(\mathbb{R}^m) \times \mathcal{P}_2(\mathbb{R}^n)$ as

$$\mathbf{F}(p, q) := \left(\frac{\exp \left\{ -(\lambda_2)^{-1} \left[\mathbb{E}_{q(y)} [f(x, y)] + \lambda_1 \|x\|^2 \right] \right\}}{Z_q}, \frac{\exp \left\{ -(\lambda_2)^{-1} \left[-\mathbb{E}_{p(x)} [f(x, y)] + \lambda_1 \|y\|^2 \right] \right\}}{Z_p} \right), \quad (35)$$

where Z_q and Z_p represent the normalization factors. Specifically,

$$Z_q = \int_{\mathbb{R}^n} \exp \left\{ -(\lambda_2)^{-1} \left[\mathbb{E}_{q(y)} [f(x, y)] + \lambda_1 \|x\|^2 \right] \right\} dx, \quad (36)$$

$$Z_p = \int_{\mathbb{R}^m} \exp \left\{ -(\lambda_2)^{-1} \left[-\mathbb{E}_{p(x)} [f(x, y)] + \lambda_1 \|y\|^2 \right] \right\} dy. \quad (37)$$

Hence, to establish the proof of the theorem, it is sufficient to demonstrate the existence of a fixed point for the operator \mathbf{F} . As established by Lemma 25, $\mathcal{P}(g, \mathcal{L}, \mathbb{R}^d)$ is a compact set for fixed g and \mathcal{L} . Therefore, the next step is to prove that for any $(p, q) \in \mathcal{P}(\mathbb{R}^m) \times \mathcal{P}(\mathbb{R}^n)$, $\mathbf{F}(p, q)$ belongs to $\mathcal{P}(g_p, \mathcal{L}_p, \mathbb{R}^m) \times \mathcal{P}(g_q, \mathcal{L}_q, \mathbb{R}^n)$ where the selection of g_p, g_q and $\mathcal{L}_p, \mathcal{L}_q$ depends on following estimation. For simplicity, we define $(\mathbf{F}_x(p, q), \mathbf{F}_y(p, q)) = \mathbf{F}(p, q)$. The formulation of Eq. (35) implies that

$$\begin{aligned} \nabla_x \mathbf{F}_x(p, q) &= -(\lambda_2)^{-1} \left[\mathbb{E}_{q(y)} [\nabla_x f(x, y)] + 2\lambda_1 x \right] \mathbf{F}_x(p, q), \\ \nabla_y \mathbf{F}_y(p, q) &= -(\lambda_2)^{-1} \left[-\mathbb{E}_{p(x)} [\nabla_y f(x, y)] + 2\lambda_1 y \right] \mathbf{F}_y(p, q). \end{aligned} \quad (38)$$

According to Assumption 1 and Proposition 26, we have

$$\begin{aligned} F_x(p, q) &\leq Z_q^{-1} \exp \left\{ (\lambda_2)^{-1} [1 - \lambda_1 \|x\|^2] \right\} \leq \gamma_m \exp \left\{ -(\lambda_2)^{-1} \lambda_1 \|x\|^2 \right\}, \\ F_y(p, q) &\leq Z_p^{-1} \exp \left\{ (\lambda_2)^{-1} [1 - \lambda_1 \|y\|^2] \right\} \leq \gamma_n \exp \left\{ -(\lambda_2)^{-1} \lambda_1 \|y\|^2 \right\}, \end{aligned}$$

and

$$\begin{aligned} \max_{x \in \mathbb{R}^m} \|\nabla_x \mathbf{F}_x(p, q)\| &\leq \gamma_m \left[(\lambda_2)^{-1} L_0 + (2\lambda_1)^{1/2} (\lambda_2)^{-1/2} \right], \\ \max_{y \in \mathbb{R}^n} \|\nabla_y \mathbf{F}_y(p, q)\| &\leq \gamma_n \left[(\lambda_2)^{-1} L_0 + (2\lambda_1)^{1/2} (\lambda_2)^{-1/2} \right], \end{aligned}$$

where $\gamma_m = \left(\frac{\lambda_1}{\lambda_2} \right)^{m/2} \pi^{-m/2} \exp \left\{ 2(\lambda_2)^{-1} \right\}$, $\gamma_n = \left(\frac{\lambda_1}{\lambda_2} \right)^{n/2} \pi^{-n/2} \exp \left\{ 2(\lambda_2)^{-1} \right\}$. Hence, we can choose

$$\begin{aligned} g_p(x) &:= \gamma_m \exp \left\{ -(\lambda_2)^{-1} \lambda_1 \|x\|^2 \right\}, \\ g_q(y) &:= \gamma_n \exp \left\{ -(\lambda_2)^{-1} \lambda_1 \|y\|^2 \right\}, \end{aligned}$$

and

$$\begin{aligned} L_{p,k} &:= \gamma_m \left[(\lambda_2)^{-1} L_0 + (2\lambda_1)^{1/2} (\lambda_2)^{-1/2} \right], \\ L_{q,k} &:= \gamma_n \left[(\lambda_2)^{-1} L_0 + (2\lambda_1)^{1/2} (\lambda_2)^{-1/2} \right], \end{aligned}$$

for any $k \in \mathbb{N}^+$. We have completed the proof w.r.t that $\mathbf{F}(p, q)$ is a continuous mapping (refer to Proposition 27) from the convex compact set $\mathcal{P}(g_p, \mathcal{L}_p, \mathbb{R}^m) \times \mathcal{P}(g_q, \mathcal{L}_q, \mathbb{R}^n)$ onto itself. Noticing that $(\mathcal{N}(0, (2\lambda_1)^{-1} \lambda_2 \mathbf{I}_m), \mathcal{N}(0, (2\lambda_1)^{-1} \lambda_2 \mathbf{I}_n))$ belongs to $\mathcal{P}(g_p, \mathcal{L}_p, \mathbb{R}^m) \times \mathcal{P}(g_q, \mathcal{L}_q, \mathbb{R}^n)$, then we finish the proof by combining Brouwer fixed-point theorem. \blacksquare

Proposition 26 aims to demonstrate that the probability density function $\mathbf{F}_x(p, q)$ and $\mathbf{F}_y(p, q)$ fall within functions family $\mathcal{B}(\mathcal{L})$ which has been mentioned in Lemma 25. Proposition 27 is proposed to establish the continuity of the operator $\mathbf{F}(p, q)$.

Proposition 26. *We have*

$$\int_{\mathbb{R}^d} \exp\{-\lambda \|x\|^2\} dx = (\lambda)^{-d/2} \pi^{d/2}, \quad (39)$$

$$\max_{x \in \mathbb{R}^d} \|x\| \exp\{-\lambda \|x\|^2\} \leq (2\lambda)^{-1/2} \exp\{-1/2\}. \quad (40)$$

Proof Considering $z \in \mathbb{R}^2$, we obtain

$$\begin{aligned} \int_{\mathbb{R}^2} \exp\{-\lambda \|z\|^2\} dz &= 2\pi \int_0^\infty r \exp\{-\lambda r^2\} dr \\ &= (\lambda)^{-1} \pi. \end{aligned} \quad (41)$$

Therefore, we derive Eq. (39) by using the fact that

$$\int_{\mathbb{R}^d} \exp\{-\lambda \|x\|^2\} dx = \left[\int_{\mathbb{R}^2} \exp\{-\lambda \|z\|^2\} dz \right]^{d/2}. \quad (42)$$

Notice that function $g(r) = r \exp\{-\lambda r^2\}$ for $r \geq 0$ attains maximum when $r = (2\lambda)^{-1/2}$. Hence, Eq. (40) is deduced from $g(r) \leq (2\lambda)^{-1/2} \exp\{-1/2\}$ directly. \blacksquare

Proposition 27. *Assuming Assumption [A₁] holds, the operator $\mathbf{F}(p, q)$ is L^1 continuous w.r.t $(p, q) \in \mathcal{P}(\mathbb{R}^m) \times \mathcal{P}(\mathbb{R}^n)$ under $\|\cdot\|_{L^1}$.*

Proof Without loss of generality, it is sufficient to provide the continuity proof for one side of p . For simplicity, we denote $\hat{\mathbf{F}}_p(y) = \exp\{-(\lambda_2)^{-1} [-\mathbb{E}_{p(x)}[f(x, y)] + \lambda_1 \|y\|^2]\}$ and $\zeta(\lambda) = \exp\{3\lambda^{-1}\} \lambda^{-1}$. According to

$$\begin{aligned} \|\hat{\mathbf{F}}_p - \hat{\mathbf{F}}_{p'}\|_{L^1} &\stackrel{a}{\leq} \int_{\mathbb{R}^n} \zeta(\lambda_2) |\mathbb{E}_{p(x)}[f(x, y)] - \mathbb{E}_{p'(x)}[f(x, y)]| \exp\{-(\lambda_2)^{-1} \lambda_1 \|y\|^2\} dy \\ &\leq \zeta(\lambda_2) \|p - p'\|_{L^1} \int_{\mathbb{R}^n} \exp\{-(\lambda_2)^{-1} \lambda_1 \|y\|^2\} dy, \end{aligned} \quad (43)$$

where (a) follows from the fact that $|\exp\{x\} - \exp\{y\}| \leq \exp\{3c\}|x - y|$ for any $x, y \in [-c, c]$ and $c > 0$, we have

$$|Z_p - Z_{p'}| \leq \|\hat{\mathbf{F}}_p - \hat{\mathbf{F}}_{p'}\|_{L^1} \leq \zeta(\lambda_2) \|p - p'\|_{L^1} \int_{\mathbb{R}^n} \exp\{-(\lambda_2)^{-1} \lambda_1 \|y\|^2\} dy. \quad (44)$$

Combining Eq. (43) and Eq. (44), we obtain

$$\begin{aligned} \left\| (Z_p)^{-1} \hat{\mathbf{F}}_p - (Z_{p'})^{-1} \hat{\mathbf{F}}_{p'} \right\|_{L^1} &\leq (Z_{p'})^{-1} \left[|Z_p - Z_{p'}| + \|\hat{\mathbf{F}}_p - \hat{\mathbf{F}}_{p'}\|_{L^1} \right] \\ &\leq \zeta(\lambda_2/2) \|p - p'\|_{L^1}. \end{aligned} \quad (45)$$

■

Appendix C. Outer Loop Error Bound

For simplicity, we define $w := (p, q)$, $\bar{w} := (\bar{p}, \bar{q})$, and

$$\begin{aligned} Q(\bar{w}, w) &:= \lambda_1 \mathbb{E}_{\bar{p}}[\|x\|^2] + \lambda_2 \mathcal{E}(\bar{p}) + \mathbb{E}_{(\bar{p}, \bar{q})}[f(x, y)] - \lambda_1 \mathbb{E}_q[\|y\|^2] - \lambda_2 \mathcal{E}(q) \\ &\quad - [\lambda_1 \mathbb{E}_p[\|x\|^2] + \lambda_2 \mathcal{E}(p) + \mathbb{E}_{(p, \bar{q})}[f(x, y)] - \lambda_1 \mathbb{E}_{\bar{q}}[\|y\|^2] - \lambda_2 \mathcal{E}(\bar{q})], \end{aligned} \quad (46)$$

$$\mathcal{Q}_x(C) := \left\{ p(x) \in \mathcal{P}_2(\mathbb{R}^m) \mid p(x) \propto e^{-l_1(x) - \frac{\lambda_1}{\lambda_2} \|x\|^2}; \|l_1\|_\infty \leq \frac{C}{\lambda_2} \right\}, \quad (47)$$

$$\mathcal{Q}_y(C) := \left\{ q(y) \in \mathcal{P}_2(\mathbb{R}^n) \mid q(y) \propto e^{-l_2(y) - \frac{\lambda_1}{\lambda_2} \|y\|^2}; \|l_2\|_\infty \leq \frac{C}{\lambda_2} \right\}, \quad (48)$$

for some positive constant C . Moreover, we denote that

$$\begin{aligned} F_\psi^x(p) &:= \int_{\mathcal{X}} (\psi(x) + \lambda_1 \|x\|^2) p(x) dx + \lambda_2 \mathcal{E}(p), \\ F_\phi^y(q) &:= \int_{\mathcal{Y}} (\phi(y) + \lambda_1 \|y\|^2) q(y) dy + \lambda_2 \mathcal{E}(q), \end{aligned}$$

for any $\psi : \mathcal{X} \rightarrow \mathbb{R}, \phi : \mathcal{Y} \rightarrow \mathbb{R}$,

$$\text{BiLin}_f(p, q) := \iint_{\mathcal{X} \times \mathcal{Y}} p(x) f(x, y) q(y) dx dy, \quad (49)$$

for any $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}, p : \mathcal{X} \rightarrow \mathbb{R}$ and $q : \mathcal{Y} \rightarrow \mathbb{R}$. In addition, we introduce some elementary symbols for $p(x) \in \mathcal{P}_2(\mathcal{X})$ satisfying LSI with constant α_p and $q(y) \in \mathcal{P}_2(\mathcal{Y})$ satisfying LSI with constant α_q , and $\delta > 0$:

$$\begin{aligned} \text{Err}_p^x(\delta) &:= \left[\lambda_2 + \frac{2\lambda_1}{\alpha_p} \right] \delta + \left[30 + \frac{8\lambda_1 \sigma(p)}{\sqrt{\alpha_p}} \right] \sqrt{\delta}, \\ \text{Err}_q^y(\delta) &:= \left[\lambda_2 + \frac{2\lambda_1}{\alpha_q} \right] \delta + \left[34 + \frac{8\lambda_1 \sigma(q)}{\sqrt{\alpha_q}} \right] \sqrt{\delta}. \end{aligned}$$

We say a density function p satisfies log-Sobolev inequality with a constant α_p i.e.

$$\mathbb{E}_p[g^2 \log(g^2)] - \mathbb{E}_p[g^2] \log \mathbb{E}_p[g^2] \leq \frac{2}{\alpha_p} \mathbb{E}_p[\|\nabla g\|^2] \quad (50)$$

In addition, we define $\sigma(p) := (\mathbb{E}_p[\|\theta\|^2])^{1/2}$ for any $p \in \mathcal{P}_2(\Pi)$. Recalling that our algorithm updates p_t and q_t according to

$$q_t^* = \arg \min_q \mathbb{E}_q[\hat{\phi}_t(y)] + \lambda_1 \mathbb{E}_q[\|y\|^2] + \lambda_2 \mathcal{E}(q) + \tau_t \text{KL}(q \| q_{t-1}^*), \quad (51)$$

$$p_t^* = \arg \min_p \mathbb{E}_p[\hat{\psi}_t(x)] + \lambda_1 \mathbb{E}_p[\|x\|^2] + \lambda_2 \mathcal{E}(p) + \eta_t \text{KL}(p \| p_{t-1}^*), \quad (52)$$

where $\hat{\phi}_t(y) = -(1 + \mu_t) \mathbb{E}_{\hat{p}_{t-1}}[f(x, y)] + \mu_t \mathbb{E}_{\hat{p}_{t-2}}[f(x, y)]$, $\hat{\psi}_t(x) = \mathbb{E}_{\hat{q}_t}[f(x, y)]$, p_t and q_t are the approximate solutions of p_t^* and q_t^* . The following Lemma summarizes the property of $p_t^*(x)$ and $q_t^*(y)$ at each step- t .

Lemma 28. *Assume $\mu_t \leq 1$ for all $t \geq 0$, then $p_t^*(x)$ and $q_t^*(y)$ satisfy*

$$p_t^*(x) \in \mathcal{Q}_x(1), \quad q_t^*(y) \in \mathcal{Q}_y(3). \quad (53)$$

Proof Note that

$$\|\hat{\phi}_t\|_\infty \leq 3, \quad \|\hat{\psi}_t\|_\infty \leq 1. \quad (54)$$

When $\mu_t \leq 1$ uniformly, we have

$$\begin{aligned} h^{(t)}(y) &= \frac{1}{\lambda_2 + \tau_t} \left\{ \hat{\phi}_t(y) + \lambda_1 \|y\|^2 + \tau_t h^{(t-1)}(y) \right\}, \\ &= \frac{\lambda_2}{\lambda_2 + \tau_t} \left\{ \frac{1}{\lambda_2} \hat{\phi}_t(y) + \frac{\lambda_1}{\lambda_2} \|y\|^2 \right\} + \frac{\tau_t}{\lambda_2 + \tau_t} h^{(t-1)}(y), \end{aligned} \quad (55)$$

$$\begin{aligned} g^{(t)}(x) &= \frac{1}{\lambda_2 + \eta_t} \left\{ \hat{\psi}_t(x) + \lambda_1 \|x\|^2 + \eta_t g^{(t-1)}(x) \right\}, \\ &= \frac{\lambda_2}{\lambda_2 + \eta_t} \left\{ \frac{1}{\lambda_2} \hat{\psi}_t(x) + \frac{\lambda_1}{\lambda_2} \|x\|^2 \right\} + \frac{\eta_t}{\lambda_2 + \eta_t} g^{(t-1)}(x), \end{aligned} \quad (56)$$

$$q_t^*(y) \propto \exp\left(-h^{(t)}(y)\right), \quad p_t^*(x) \propto \exp\left(-g^{(t)}(x)\right). \quad (57)$$

Suppose that $p_{t-1}^*(x) \in \mathcal{Q}_x(1)$, $q_{t-1}^*(y) \in \mathcal{Q}_y(3)$ which indicates that $g^{t-1}(x) = -l_{1,t-1}(x) - \frac{\lambda_1}{\lambda_2} \|x\|^2$ with $\|l_{1,t-1}\|_\infty \leq \frac{1}{\lambda_2}$ and $h^{t-1}(y) = -l_{2,t-1}(y) - \frac{\lambda_1}{\lambda_2} \|y\|^2$ with $\|l_{2,t-1}\|_\infty \leq \frac{3}{\lambda_2}$, we obtain

$$\begin{aligned} \|l_{1,t}\|_\infty &\leq \frac{\lambda_2}{\lambda_2 + \eta_t} \left\| \frac{1}{\lambda_2} \hat{\psi} \right\|_\infty + \frac{\eta_t}{\lambda_2 + \eta_t} \|l_{1,t-1}\|_\infty \leq \frac{1}{\lambda_2}, \\ \|l_{2,t}\|_\infty &\leq \frac{\lambda_2}{\lambda_2 + \eta_t} \left\| \frac{1}{\lambda_2} \hat{\phi} \right\|_\infty + \frac{\eta_t}{\lambda_2 + \eta_t} \|l_{2,t-1}\|_\infty \leq \frac{1}{\lambda_2}. \end{aligned} \quad (58)$$

Hence, the desired results can be obtained by combining the fact that $p_0^*(x) \in \mathcal{Q}_x(1)$, $q_0^*(y) \in \mathcal{Q}_y(3)$ and mathematical induction. \blacksquare

We use Lemma 29 to deal with the error bound between inexact solution p_t and exact solution p_t^* (as well as q_t and q_t^*) in Theorem 31.

Lemma 29. *Assume a probability density $\bar{p}(\theta) \propto \exp(-l(\theta) - \frac{\lambda_1}{\lambda_2}\|\theta\|^2)$ with bounded $l : \mathbb{R}^{d_\theta} \rightarrow \mathbb{R}$ which satisfies $\|l\|_{L^\infty} \leq C$, then for all probability density $p(\theta) \in \mathcal{P}_2(\mathbb{R}^{d_\theta})$ we have*

$$\mathcal{E}(p) - \mathcal{E}(\bar{p}) \leq \text{KL}(p\|\bar{p}) + \left(C + \frac{4\lambda_1\sigma(\bar{p})}{\lambda_2\sqrt{\alpha_{\bar{p}}}}\right) \sqrt{\text{KL}(p\|\bar{p})}, \quad (59)$$

$$\int \|\theta\|^2(p - \bar{p})(\theta)d\theta \leq \frac{2}{\alpha_{\bar{p}}}\text{KL}(p\|\bar{p}) + \frac{4\sigma(\bar{p})}{\sqrt{\alpha_{\bar{p}}}} \sqrt{\text{KL}(p\|\bar{p})}. \quad (60)$$

Proof According to the definition of $\mathcal{E}(\cdot)$, we obtain

$$\begin{aligned} \mathcal{E}(p) - \mathcal{E}(\bar{p}) &= \text{KL}(p\|\bar{p}) + \int \log(\bar{p}(\theta))(p - \bar{p})(\theta)d\theta \\ &= \text{KL}(p\|\bar{p}) + \int \left(l(\theta) + \frac{\lambda_1}{\lambda_2}\|\theta\|^2\right) (\bar{p} - p)(\theta)d\theta \\ &\leq \text{KL}(p\|\bar{p}) + C\|\bar{p} - p\|_{\text{TV}} + \frac{\lambda_1}{\lambda_2} \int \|\theta\|^2(\bar{p} - p)(\theta)d\theta. \end{aligned} \quad (61)$$

Therefore, in order to complete the proof, we just need to bound $|\int \|\theta\|^2(p - \bar{p})(\theta)d\theta|$. Letting $\rho(\theta, \theta')$ denote an optimal coupling between $p(\theta)$ and $\bar{p}(\theta)$, we have

$$\begin{aligned} \int \|\theta\|^2(\bar{p} - p)(\theta)d\theta &= \int \|\theta'\|^2\bar{p}(\theta')d\theta' - \int \|\theta\|^2p(\theta)d\theta \\ &= \iint (\langle \theta' - \theta, \theta' + \theta \rangle) \rho(\theta, \theta')d\theta d\theta' \\ &= \iint (-\|\theta' - \theta\|^2 + 2\langle \theta' - \theta, \theta' \rangle) \rho(\theta, \theta')d\theta d\theta' \\ &\leq 2W_2(p, \bar{p})(\mathbb{E}_{\bar{p}}[\|\theta\|^2])^{1/2}, \end{aligned} \quad (62)$$

where the last inequality in Eq. (62) follows from Hölder's inequality. Then, applying Eq. (62) on Eq. (61) and combining Talagrand's inequality 20, Pinsker's inequality $\|\bar{p} - p\|_{\text{TV}} \leq \sqrt{\text{KL}(p\|\bar{p})/2}$, we conclude Eq. (59). On the other hand, we can obtain

$$\begin{aligned} \int \|\theta\|^2(p - \bar{p})(\theta)d\theta &= \int \|\theta\|^2p(\theta)d\theta - \int \|\theta'\|^2\bar{p}(\theta')d\theta' \\ &= \iint (\|\theta\|^2 - \|\theta'\|^2) \rho(\theta, \theta')d\theta d\theta' \\ &= \iint (\|\theta - \theta'\|^2 + 2\langle \theta - \theta', \theta' \rangle) \rho(\theta, \theta')d\theta d\theta' \\ &\leq W_2^2(p, \bar{p}) + 2W_2(p, \bar{p})(\mathbb{E}_{\bar{p}}[\|\theta\|^2])^{1/2}, \end{aligned} \quad (63)$$

where the last inequality in Eq. (63) follows from Hölder's inequality as well. Finally, we deduce Eq. (60) by combining Eq. (63) and Talagrand's inequality 20. \blacksquare

Lemma 3 describe the form of optimal solution in each iteration which implies that if p_{t-1}^* and q_{t-1}^* are proximal Gibbs distribution, then p_t^* and q_t^* have the similar expression as well.

Proof [Proof of Lemma 3] For simplicity, assume that

$$F(p') \stackrel{\text{def}}{=} \mathbb{E}_{p'}[l(\theta)] + \lambda \mathcal{E}(p') + \tau \text{KL}(p' \|\bar{p}), \quad (64)$$

for any $p' \in \mathcal{P}(\Pi)$.

Clearly, $p \in \mathcal{P}(\Pi)$ and $p(\theta) \propto (\bar{p}(\theta))^{\frac{\tau}{\lambda+\tau}} \exp\left\{-\frac{1}{\lambda+\tau}l(\theta)\right\}$. Hence, for any $p' \in \mathcal{P}(\Pi)$,

$$\begin{aligned} F(p') &= \mathbb{E}_{p'}[l(\theta)] + \lambda \mathcal{E}(p') + \tau \text{KL}(p' \|\bar{p}) \\ &= \mathbb{E}_p[l(\theta)] + \lambda \mathcal{E}(p) + \tau \text{KL}(p \|\bar{p}) + \int (p' - p)(\theta)[l(\theta) - \tau \log(\bar{p}(\theta))]d\theta \\ &\quad + (\lambda + \tau)[\mathcal{E}(p') - \mathcal{E}(p)] \\ &= F(p) + \int (p' - p)[-(\lambda + \tau) \log(p)](\theta)d\theta + (\lambda + \tau)[\mathcal{E}(p') - \mathcal{E}(p)] \\ &= F(p) + (\lambda + \tau) \text{KL}(p' \|\bar{p}) \geq F(p). \end{aligned} \quad (65)$$

■

By choosing $l(\theta) = f(\theta) + \lambda_1 \|\theta\|^2$ and the probability density $\bar{p}(\theta) \propto \exp(-\tilde{l}(\theta))$, we obtain that the minimizer of

$$\mathbb{E}_p[f(\theta)] + \lambda_1 \mathbb{E}_p[\|\theta\|^2] + \lambda \mathcal{E}(p) + \tau \text{KL}(p \|\bar{p}),$$

on $\mathcal{P}(\Pi)$ is $p \propto \exp(-\frac{1}{\lambda+\tau}f(\theta) - \frac{\tau}{\lambda+\tau}\tilde{l}(\theta) - \frac{\lambda_1}{\lambda+\tau}\|\theta\|^2)$.

Lemma 30 indicates the property of optimal solution in each step- t , which plays a crucial role in evaluating the outer loop error bound (in Theorem 31).

Lemma 30. *Assume that $l : \Pi \rightarrow \mathbb{R}$. If*

$$\hat{p} = \arg \min_p \{\mathbb{E}_p[l(\theta)] + \mu \mathcal{E}(p) + \text{KL}(p \|\tilde{p})\}, \quad (66)$$

then

$$\mathbb{E}_{\hat{p}}[l(\theta)] + \mu \mathcal{E}(\hat{p}) + \text{KL}(\hat{p} \|\tilde{p}) + (\mu + 1) \text{KL}(p \|\hat{p}) = \mathbb{E}_p[l(\theta)] + \mu \mathcal{E}(p) + \text{KL}(p \|\tilde{p}), \quad \forall p \in \mathcal{P}.$$

Proof The conclusion of this lemma can be derived directly from Eq. (65) by applying $\lambda = \mu$ and $\tau = 1$. ■

Theorem 31 (Long Version of Theorem 8). *Assuming Assumptions $[A_1]$, $[A_2]$ and $[A_3]$ hold, we set*

$$\gamma_t \mu_t = \gamma_{t-1}, \quad (67)$$

$$\gamma_t \tau_t \leq \gamma_{t-1}(\tau_{t-1} + \lambda_2), \quad (68)$$

$$\gamma_t \eta_t \leq \gamma_{t-1}(\eta_{t-1} + \lambda_2), \quad (69)$$

$$\tau_t \eta_{t-1} \geq \mu_t, \quad (70)$$

$$\mu_t \leq 1, \quad (71)$$

and assume that subproblem error bounds at each step- t satisfy

$$\text{KL}(p_t \| p_t^*) \leq \delta_{t,1}, \text{KL}(q_t \| q_t^*) \leq \delta_{t,2}, \quad (72)$$

then

$$\begin{aligned} \sum_{t=1}^k \gamma_t Q(w_t, w) &\leq \gamma_1 \eta_1 \text{KL}(p \| p_0^*) - \gamma_k (\eta_k + \lambda_2) \text{KL}(p \| p_k^*) + \gamma_1 \tau_1 \text{KL}(q \| q_0^*) \\ &\quad - \gamma_k \left(\tau_k + \lambda_2 - \frac{1}{4\eta_k} \right) \text{KL}(q \| q_k^*) \\ &\quad + \sum_{t=1}^k \gamma_t \left[4(1 + \mu_t) \epsilon + \text{Err}_{p_t^*}^x(\delta_{t,1}) + \text{Err}_{q_t^*}(\delta_{t,2}) \right], \end{aligned} \quad (73)$$

for any $\epsilon > 0$ and $p \in \mathcal{Q}_x(C_1), q \in \mathcal{Q}_y(C_2)$ with probability at least $1 - \delta$ for sample size $M = \Theta(\epsilon^{-2} [\log(T\delta^{-1}) + d \log(1 + L_0 r_{\max\{C_1, C_2\}}(\epsilon, \lambda_1, \lambda_2) \epsilon^{-1})])$ in each iteration where $d = \max\{m, n\}$ and $r_C(\epsilon, \lambda_1, \lambda_2) := 8\sqrt{\frac{\lambda_2 + C}{\lambda_1} \max\{\log(\epsilon^{-1}), m, n\}}$. Moreover, we have

$$\begin{aligned} \gamma_k \left(\tau_k + \lambda_2 - \frac{1}{4\eta_k} \right) \text{KL}(q_* \| q_k^*) &\leq \gamma_1 \eta_1 \text{KL}(p_* \| p_0^*) + \gamma_1 \tau_1 \text{KL}(q_* \| q_0^*) \\ &\quad + \sum_{t=1}^k \gamma_t \left[4(1 + \mu_t) \epsilon + \text{Err}_{p_t^*}(\delta_{t,1}) + \text{Err}_{q_t^*}(\delta_{t,2}) \right], \end{aligned} \quad (74)$$

and

$$\begin{aligned} \gamma_k \left(\eta_k - \frac{1}{4(\tau_k + \lambda_2)} \right) \text{KL}(p_k^* \| p_{k-1}^*) &\leq \gamma_1 \eta_1 \text{KL}(p_* \| p_0^*) + \gamma_1 \tau_1 \text{KL}(q_* \| q_0^*) \\ &\quad + \sum_{t=1}^k \gamma_t \left[4(1 + \mu_t) \epsilon + \text{Err}_{p_t^*}(\delta_{t,1}) + \text{Err}_{q_t^*}(\delta_{t,2}) \right], \end{aligned} \quad (75)$$

with respect to the MNE $w_* = (p_*, q_*)$ of Eq. (2).

Proof By applying Lemma 30 to Eq. (51) and Eq. (52), we can obtain the following equalities

$$F_{\hat{\phi}_t}^y(q_t^*) - F_{\hat{\phi}_t}^y(q) = \tau_t [\text{KL}(q \| q_{t-1}^*) - \text{KL}(q_t^* \| q_{t-1}^*)] - (\tau_t + \lambda_2) \text{KL}(q \| q_t^*), \quad (76)$$

$$F_{\hat{\psi}_t}^x(p_t^*) - F_{\hat{\psi}_t}^x(p) = \eta_t [\text{KL}(p \| p_{t-1}^*) - \text{KL}(p_t^* \| p_{t-1}^*)] - (\eta_t + \lambda_2) \text{KL}(p \| p_t^*). \quad (77)$$

According to Eq. (76), we have

$$\begin{aligned}
 F_{\hat{\phi}_t}^y(q_t) - F_{\hat{\phi}_t}^y(q) &= F_{\hat{\phi}_t}^y(q_t^*) - F_{\hat{\phi}_t}^y(q) - \left[F_{\hat{\phi}_t}^y(q_t^*) - F_{\hat{\phi}_t}^y(q_t) \right] \\
 &= \tau_t [\text{KL}(q \| q_{t-1}^*) - \text{KL}(q_t^* \| q_{t-1}^*)] - (\tau_t + \lambda_2) \text{KL}(q \| q_t^*) \\
 &\quad + \underbrace{\lambda_1 \int \|y\|^2 (q_t - q_t^*)(y) dy + \lambda_2 (\mathcal{E}(q_t) - \mathcal{E}(q_t^*))}_a \\
 &\quad + \underbrace{\int \hat{\phi}_t(y) (q_t - q_t^*)(y) dy}_b.
 \end{aligned} \tag{78}$$

It is worth noting that we can derive the following inequality from Eq. (71)

$$\int \hat{\phi}_t(y) (q_t - q_t^*)(y) dy \leq 3 \|q_t^* - q_t\|_{\text{TV}} \leq 3 \sqrt{\text{KL}(q_t \| q_t^*)}. \tag{79}$$

Next, we apply Lemma 29 to (a) and bound Eq. (79) in (b) of Eq. (78). This allows us to propose the estimation as follows

$$\begin{aligned}
 F_{\hat{\phi}_t}^y(q_t) - F_{\hat{\phi}_t}^y(q) &\leq \tau_t [\text{KL}(q \| q_{t-1}^*) - \text{KL}(q_t^* \| q_{t-1}^*)] - (\tau_t + \lambda_2) \text{KL}(q \| q_t^*) \\
 &\quad + \left[\lambda_2 + \frac{2\lambda_1}{\alpha_{q_t^*}} \right] \delta_{t,2} + \left[6 + \frac{8\lambda_1 \sigma(q_t^*)}{\sqrt{\alpha_{q_t^*}}} \right] \sqrt{\delta_{t,2}}.
 \end{aligned} \tag{80}$$

Likewise, we can deduce that

$$\begin{aligned}
 F_{\hat{\psi}_t}^x(p_t^*) - F_{\hat{\psi}_t}^x(p) &\leq \eta_t [\text{KL}(p \| p_{t-1}^*) - \text{KL}(p_t^* \| p_{t-1}^*)] - (\eta_t + \lambda_2) \text{KL}(p \| p_t^*) \\
 &\quad + \left[\lambda_2 + \frac{2\lambda_1}{\alpha_{p_t^*}} \right] \delta_{t,1} + \left[2 + \frac{8\lambda_1 \sigma(p_t^*)}{\sqrt{\alpha_{p_t^*}}} \right] \sqrt{\delta_{t,1}}.
 \end{aligned} \tag{81}$$

By utilizing the boundedness of f and carefully selecting constants C_1, C_2 , we can adjust the number of sampling points for both \hat{p}_t and \hat{q}_t to obtain the following result

$$|\text{BiLin}_f(p_t - \hat{p}_t, q)| \leq \epsilon, \quad |\text{BiLin}_f(p, q_t - \hat{q}_t)| \leq \epsilon, \tag{82}$$

with probability $1 - \delta$ for all $q(y) \in \mathcal{Q}_Y(C_2)$, $p(x) \in \mathcal{Q}_X(C_1)$ uniformly (Refer to Lemma 7 in Appendix F). In order to deal with finite particles, we denote $\phi_t(y) = -(1 + \mu_t) \mathbb{E}_{p_{t-1}}[f(x, y)] + \mu_t \mathbb{E}_{p_{t-2}}[f(x, y)]$, $\psi_t(x) = \mathbb{E}_{q_t}[f(x, y)]$. Therefore, summing up these inequalities in Eq. (80),

Eq. (81) and using the definition of the Max-Min gap function Q , we have

$$\begin{aligned}
 & Q(w_t, w) + \text{BiLin}_f((p_t - p_{t-1}) - \mu_t(p_{t-1} - p_{t-2}), q_t - q) \\
 & \stackrel{\text{c}}{=} F_{\hat{\phi}_t}^y(q_t) - F_{\hat{\phi}_t}^y(q) + F_{\hat{\psi}_t}^x(p_t) - F_{\hat{\psi}_t}^x(p) \\
 & = \underbrace{\int (\phi_t - \hat{\phi}_t)(y)(q_t^* - q)(y)dy}_{\text{bounded by Eq. (82)}} + \underbrace{\int (\phi_t - \hat{\phi}_t)(y)(q_t - q_t^*)(y)dy}_{\text{bounded by Pinsker's inequality}} \\
 & \quad + \underbrace{\int (\psi_t - \hat{\psi}_t)(x)(p_t^* - p)(x)dx}_{\text{bounded by Eq. (82)}} + \underbrace{\int (\psi_t - \hat{\psi}_t)(x)(p_t - p_t^*)(x)dx}_{\text{bounded by Pinsker's inequality}} \\
 & \quad + F_{\hat{\phi}_t}^y(q_t) - F_{\hat{\phi}_t}^y(q) + F_{\hat{\psi}_t}^x(p_t) - F_{\hat{\psi}_t}^x(p) \\
 & \stackrel{\text{d}}{\leq} 4(1 + \mu_t)\epsilon + \tau_t[\text{KL}(q||q_{t-1}^*) - \text{KL}(q_t^*||q_{t-1}^*)] - (\tau_t + \lambda_2)\text{KL}(q||q_t^*) \\
 & \quad + \eta_t[\text{KL}(p||p_{t-1}^*) - \text{KL}(p_t^*||p_{t-1}^*)] - (\eta_t + \lambda_2)\text{KL}(p||p_t^*) + \left[\lambda_2 + \frac{2\lambda_1}{\alpha_{p_t^*}} \right] \delta_{t,1} \\
 & \quad + \left[6 + \frac{8\lambda_1\sigma(p_t^*)}{\sqrt{\alpha_{p_t^*}}} \right] \sqrt{\delta_{t,1}} + \left[\lambda_2 + \frac{2\lambda_1}{\alpha_{q_t^*}} \right] \delta_{t,2} + \left[18 + \frac{8\lambda_1\sigma(q_t^*)}{\sqrt{\alpha_{q_t^*}}} \right] \sqrt{\delta_{t,2}}, \tag{83}
 \end{aligned}$$

where equality (c) follows from the expression of Q and inequality (d) is deduced from plugging Eq. (80) into Eq. (81).

Notice that

$$\begin{aligned}
 & \text{BiLin}_f((p_t - p_{t-1}) - \mu_t(p_{t-1} - p_{t-2}), q_t - q) \\
 & = \text{BiLin}_f(p_t - p_{t-1}, q_t - q) - \mu_t \text{BiLin}_f(p_{t-1} - p_{t-2}, q_{t-1} - q) \\
 & \quad + \mu_t \text{BiLin}_f(p_{t-1} - p_{t-2}, q_{t-1} - q_t). \tag{84}
 \end{aligned}$$

Moreover, we can derive

$$\begin{aligned}
 & \text{BiLin}_f((p_t^* - p_{t-1}^*) - \mu_t(p_{t-1}^* - p_{t-2}^*), q_t^* - q) \\
 & \quad - \text{BiLin}_f((p_t - p_{t-1}) - \mu_t(p_{t-1} - p_{t-2}), q_t - q) \\
 & \leq 8 \left[\sqrt{\delta_{t,1}} + \sqrt{\delta_{t-1,1}} + \sqrt{\delta_{t-2,1}} + \sqrt{\delta_{t,2}} \right], \tag{85}
 \end{aligned}$$

where Eq. (85) follows from

$$\begin{aligned}
 & \text{BiLin}_f((p_t^* - p_{t-1}^*) - \mu_t(p_{t-1}^* - p_{t-2}^*), q_t^* - q) \\
 & \quad - \text{BiLin}_f((p_t - p_{t-1}) - \mu_t(p_{t-1} - p_{t-2}), q_t - q) \\
 & = \text{BiLin}_f((p_t^* - p_t) - (1 + \mu_t)(p_{t-1}^* - p_{t-1}) + \mu_t(p_{t-2}^* - p_{t-2}), q_t^* - q) \\
 & \quad + \text{BiLin}_f((p_t - p_{t-1}) - \mu_t(p_{t-1} - p_{t-2}), q_t^* - q_t) \\
 & \stackrel{\text{e}}{\leq} 2[\|p_t^* - p_t\|_{\text{TV}} + (1 + \mu_t)\|p_{t-1}^* - p_{t-1}\|_{\text{TV}} \\
 & \quad + \mu_t\|p_{t-2}^* - p_{t-2}\|_{\text{TV}} + (1 + \mu_t)\|q_t^* - q_t\|_{\text{TV}}], \tag{86}
 \end{aligned}$$

and Inequality (e) can be deduced from the boundness of $\int f(x, y)(q_t^* - q)(y)dy$, $\int f(x, y)(p_t - p_{t-1})(x)dx$ and $\int f(x, y)(p_{t-1} - p_{t-2})(x)dx$.

By applying $p_{t-2}^*, p_{t-1}^*, p_t^*, q_{t-1}^*, q_t^*$ to Eq. (84), and then combining the two aforementioned relations Eq. (83) and Eq. (85), we can multiply both sides of the resulting inequality by $\gamma_t \geq 0$ to obtain

$$\begin{aligned}
 \sum_{t=1}^k \gamma_t Q(w_t, w) &\leq \sum_{t=1}^k \gamma_t [\eta_t \text{KL}(p \| p_{t-1}^*) - (\eta_t + \lambda_2) \text{KL}(p \| p_t^*)] \\
 &\quad + \sum_{t=1}^k \gamma_t [\tau_t \text{KL}(q \| q_{t-1}^*) - (\tau_t + \lambda_2) \text{KL}(q \| q_t^*)] + 4 \sum_{t=1}^k \gamma_t (1 + \mu_t) \epsilon \\
 &\quad + \sum_{t=1}^k \gamma_t [\mu_t \text{BiLin}_f(p_{t-1}^* - p_{t-2}^*, q_{t-1}^* - q) - \text{BiLin}_f(p_t^* - p_{t-1}^*, q_t^* - q)] \\
 &\quad - \sum_{t=1}^k \gamma_t [\tau_t \text{KL}(q_t^* \| q_{t-1}^*) + \eta_t \text{KL}(p_t^* \| p_{t-1}^*) + \mu_t \text{BiLin}_f(p_{t-1}^* - p_{t-2}^*, q_{t-1}^* - q_t^*)] \\
 &\quad + \sum_{t=1}^k \gamma_t [\text{Err}_{p_t^*}^x(\delta_{t,1}) + \text{Err}_{q_t^*}^y(\delta_{t,2})]. \tag{87}
 \end{aligned}$$

Taking into account Eq. (67)-(69) and the fact that $p_0^* = p_{-1}^*$ and $q_0^* = q_{-1}^*$, the above inequality Eq. (87) implies that

$$\begin{aligned}
 \sum_{t=1}^k \gamma_t Q(w_t, w) &\leq \gamma_1 \eta_1 \text{KL}(p \| p_0^*) - \gamma_k (\eta_k + \lambda_2) \text{KL}(p \| p_k^*) \\
 &\quad + \gamma_1 \tau_1 \text{KL}(q \| q_0^*) - \gamma_k (\tau_k + \lambda_2) \text{KL}(q \| q_k^*) + 4 \sum_{t=1}^k \gamma_t (1 + \mu_t) \epsilon \\
 &\quad - \sum_{t=1}^k \gamma_t [\tau_t \text{KL}(q_t^* \| q_{t-1}^*) + \eta_t \text{KL}(p_t^* \| p_{t-1}^*) + \mu_t \text{BiLin}_f(p_{t-1}^* - p_{t-2}^*, q_{t-1}^* - q_t^*)] \\
 &\quad - \gamma_k \text{BiLin}_f(p_k^* - p_{k-1}^*, q_k^* - q) + \sum_{t=1}^k \gamma_t [\text{Err}_{p_t^*}^x(\delta_{t,1}) + \text{Err}_{q_t^*}^y(\delta_{t,2})]. \tag{88}
 \end{aligned}$$

By Eq. (67) and Eq. (70), we have

$$\begin{aligned}
 &- \sum_{t=1}^k \gamma_t [\tau_t \text{KL}(q_t^* \| q_{t-1}^*) + \eta_t \text{KL}(p_t^* \| p_{t-1}^*) + \mu_t \text{BiLin}_f(p_{t-1}^* - p_{t-2}^*, q_{t-1}^* - q_t^*)] \\
 \stackrel{\text{f}}{\leq} &- \sum_{t=2}^k [\gamma_t \tau_t \text{KL}(q_t^* \| q_{t-1}^*) + \gamma_{t-1} \eta_{t-1} \text{KL}(p_{t-1}^* \| p_{t-2}^*) \\
 &\quad - \gamma_t \mu_t \|p_{t-1}^* - p_{t-2}^*\|_{\text{TV}} \|q_{t-1}^* - q_t^*\|_{\text{TV}}] - \gamma_k \eta_k \text{KL}(p_k^* \| p_{k-1}^*) \\
 \stackrel{\text{g}}{\leq} &- \gamma_k \eta_k \text{KL}(p_k^* \| p_{k-1}^*), \tag{89}
 \end{aligned}$$

where the inequality (f) follows from

$$\text{BiLin}_f(p_{t-1}^* - p_{t-2}^*, q_{t-1}^* - q_t^*) \leq \|p_{t-1}^* - p_{t-2}^*\|_{\text{TV}} \|q_{t-1}^* - q_t^*\|_{\text{TV}}, \tag{90}$$

and the inequality (g) is derived from the Pinsker's inequality and the parameter settings in Eq. (67) and Eq. (70).

Combining the two inequalities Eq. (88) and Eq. (89), we obtain

$$\begin{aligned}
 \sum_{t=1}^k \gamma_t Q(w_t, w) &\leq \gamma_1 \eta_1 \text{KL}(p \| p_0^*) - \gamma_k (\eta_k + \lambda_2) \text{KL}(p \| p_k^*) + \gamma_1 \tau_1 \text{KL}(q \| q_0^*) \\
 &\quad - \gamma_k (\tau_k + \lambda_2) \text{KL}(q \| q_k^*) - \gamma_k \eta_k \text{KL}(p_k^* \| p_{k-1}^*) - \gamma_k \text{BiLin}_f(p_k^* - p_{k-1}^*, q_k^* - q) \\
 &\quad + 4 \sum_{t=1}^k \gamma_t (1 + \mu_t) \epsilon + \sum_{t=1}^k \gamma_t \left[\text{Err}_{p_t^*}^x(\delta_{t,1}) + \text{Err}_{q_t^*}^y(\delta_{t,2}) \right]. \tag{91}
 \end{aligned}$$

The result in Eq. (73) then follows from the above inequality (91) and the fact that by Eq. (70),

$$\begin{aligned}
 & - (\tau_k + \lambda_2) \text{KL}(q \| q_k^*) - \eta_k \text{KL}(p_k^* \| p_{k-1}^*) - \text{BiLin}_f(p_k^* - p_{k-1}^*, q_k^* - q) \\
 & \leq - (\tau_k + \lambda_2) \text{KL}(q \| q_k^*) - \eta_k \text{KL}(p_k^* \| p_{k-1}^*) + \sqrt{\text{KL}(p_k^* \| p_{k-1}^*)} \sqrt{\text{KL}(q \| q_k^*)} \\
 & \leq - \left(\tau_k + \lambda_2 - \frac{1}{4\eta_k} \right) \text{KL}(q \| q_k^*). \tag{92}
 \end{aligned}$$

Fixing $w = w_*$ in the above inequality (91) and combining Eq. (92), we obtain the result in Eq. (74).

Finally, Eq. (75) is concluded from similar ideas,

$$\begin{aligned}
 & - (\tau_k + \lambda_2) \text{KL}(q \| q_k^*) - \eta_k \text{KL}(p_k^* \| p_{k-1}^*) - \text{BiLin}_f(p_k^* - p_{k-1}^*, q_k^* - q) \\
 & \leq - (\tau_k + \lambda_2) \text{KL}(q \| q_k^*) - \eta_k \text{KL}(p_k^* \| p_{k-1}^*) + \sqrt{\text{KL}(p_k^* \| p_{k-1}^*)} \sqrt{\text{KL}(q \| q_k^*)} \\
 & \leq - \left(\eta_k - \frac{1}{4(\tau_k + \lambda_2)} \right) \text{KL}(p_k^* \| p_{k-1}^*). \tag{93}
 \end{aligned}$$

■

Based on Theorem 31, we will provide specific instantiation methods for parameters $\{\tau_t\}_{t=1}^k, \{\eta_t\}_{t=1}^k, \{\gamma_t\}_{t=1}^k$. Additionally, leveraging Theorem 11, we can control the error bounds $\{\delta_{t,1}\}_{t=1}^k, \{\delta_{t,2}\}_{t=1}^k$ (refer to Theorem 31) of sub-problems. By merging these two aspects, we can achieve a globally convergent result for solving the minimax problem with accelerated convergence rates. Our goal is to establish the convergence of the Max-Min gap, as evaluated at the output solution

$$\bar{w}_k = \frac{\sum_{t=1}^k \gamma_t w_t}{\sum_{t=1}^k \gamma_t}, \tag{94}$$

will ultimately converge to zero. In addition, We will establish the last iterate convergence of the PAPAL algorithm by considering the KL distance between w_k^* and w_* in KL divergence, as well as the W_2 distance between w_k and w_* .

For Corollary 9, we have the KL distance estimation:

$$\begin{aligned} \text{KL}(p_* \| p_T^*) + \frac{1}{2} \text{KL}(q_* \| q_T^*) &\leq \mu^T [\text{KL}(q_* \| q_0^*) + \text{KL}(p_* \| p_0^*)] \\ &\quad + \lambda_2^{-1} [8\epsilon + (10\lambda_1 + \lambda_2 + 32)T^{-J}], \end{aligned} \quad (95)$$

and the Wasserstein distance estimation:

$$\begin{aligned} \frac{1}{2} W_2^2(q_T, q_*) + W_2^2(p_T, p_*) &\leq \frac{4}{\min\{\alpha_{p_T^*}, \alpha_{q_T^*}\}} \{ \mu^T [\text{KL}(q_* \| q_0^*) + \text{KL}(p_* \| p_0^*)] \\ &\quad + \lambda_2^{-1} [8\epsilon + (10\lambda_1 + \lambda_2 + 32)T^{-J}] \} + 3T^{-2J}, \end{aligned} \quad (96)$$

with respect to the MNE $w_* = (p_*, q_*)$ of Eq. (2).

Proof [Proof of Corollary 9] According to the sub-problem error bound, we have

$$\frac{2\lambda_1}{\alpha_{p_t^*}} \delta_{t,1} \leq \lambda_1 T^{-2J}, \quad \frac{8\lambda_1 \sigma(p_t^*)}{\sqrt{\alpha_{p_t^*}}} \sqrt{\delta_{t,1}} \leq 4\lambda_1 T^{-J}, \quad \frac{2\lambda_1}{\alpha_{q_t^*}} \delta_{t,2} \leq \lambda_1 T^{-2J}, \quad \frac{8\lambda_1 \sigma(q_t^*)}{\sqrt{\alpha_{q_t^*}}} \sqrt{\delta_{t,2}} \leq 4\lambda_1 T^{-J}.$$

Notice that $Q(\bar{w}, w)$ is convex with respect to \bar{w} given a fixing w .

According to Jensen's inequality, we have

$$Q(\bar{w}_T, w) \leq \frac{\sum_{t=1}^T \gamma_t Q(w_t, w)}{\sum_{t=1}^T \gamma_t}, \quad (97)$$

for any $w \in \mathcal{P}_2(\mathcal{X}) \times \mathcal{P}_2(\mathcal{Y})$. Moreover, under problem setting in Eq. 2 and the optimality condition, we can derive that $p_* \in \mathcal{Q}_x(C)$, $q_* \in \mathcal{Q}_y(C)$ when w_* exists, and $\bar{p}^* \in \mathcal{Q}_x(C)$, $\bar{q}^* \in \mathcal{Q}_y(C)$ for any $\bar{w} = (\bar{p}, \bar{q}) \in \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y})$ with $C = 1$ (see coefficient C in (47) and (48)).

Therefore, Eq. (25) is derived from Eq. (73)(with $w = \bar{w}_T^*$),

$$\begin{aligned} 0 \leq Q(\bar{w}_T, \bar{w}_T^*) &\leq \frac{\sum_{t=1}^T \gamma_t Q(w_T, \bar{w}_T^*)}{\sum_{t=1}^T \gamma_t} \leq \mu^T [\gamma_1 \eta_1 \text{KL}(p \| p_0^*) + \gamma_1 \tau_1 \text{KL}(q \| q_0^*)] \\ &\quad + \frac{\sum_{t=1}^T \gamma_t}{\sum_{t=1}^T \gamma_t} [8\epsilon + (10\lambda_1 + \lambda_2 + 32)T^{-J}]. \end{aligned} \quad (98)$$

In addition, we have

$$2\eta_t(\tau_t + \lambda_2) - 1 = 2\lambda_2^2 \frac{\mu}{(1-\mu)^2} = 1,$$

which leads to

$$\frac{\tau_t + \lambda_2}{2} - \frac{1}{4\eta_t} > 0. \quad (99)$$

The result in Eq. (22) then follows from Eq. (73) (with $w = w_*$) and $Q(w_t, w_*) \geq 0$ that

$$\begin{aligned} 0 \leq \sum_{t=1}^T \gamma_t Q(w_t, w_*) &\leq \gamma_1 \eta_1 \text{KL}(p_* \| p_0^*) - \gamma_T (\eta_T + \lambda_2) \text{KL}(p_* \| p_T^*) + \gamma_1 \tau_1 \text{KL}(q_* \| q_0^*) \\ &\quad - \frac{\gamma_T}{2} (\tau_T + \lambda_2) \text{KL}(q_* \| q_T^*) \\ &\quad + \sum_{t=1}^T \gamma^t [8\epsilon + (10\lambda_1 + \lambda_2 + 32)T^{-J}]. \end{aligned} \quad (100)$$

Finally, by combining Eq. (22), Talagrand's inequality, and

$$W_2^2(\nu, \rho) \leq 2W_2^2(\nu, p) + 2W_2^2(p, \rho), \quad \forall \nu, p, \rho \in \mathcal{P}_2(\Pi), \quad (101)$$

we obtain the convergence of Wasserstein distance. \blacksquare

Appendix D. Global Convergence

Proof [Proof of Theorem 12] Noting the explicit form of $p_t^*(x)$ and $q_t^*(y)$ in Lemma 28, we obtain that p_t^* satisfies log-Sobolev inequality with $\lambda_1 \exp\{-4\lambda_2^{-1}\}\lambda_2^{-1}$ and q_t^* satisfies log-Sobolev inequality with $\lambda_1 \exp\{-12\lambda_2^{-1}\}\lambda_2^{-1}$ for any $t \geq 0$ by combining Lemma 10.

In fact, we also have

$$\begin{aligned} \mathbb{E}_{\bar{p}}[\|\theta\|^2] &= \int \|\theta\|^2 \frac{\exp[-l(\theta) - (\lambda_1/\lambda_2)\|\theta\|^2]}{\int \exp[-l(\theta) - (\lambda_1/\lambda_2)\|\theta\|^2]} d\theta \\ &\leq \exp(2C) \int \|\theta\|^2 \frac{\exp[-(\lambda_1/\lambda_2)\|\theta\|^2]}{\int \exp[-(\lambda_1/\lambda_2)\|\theta\|^2]} d\theta \\ &= \frac{\lambda_2 d_\theta}{2\lambda_1} \exp(2C), \end{aligned} \quad (102)$$

if probability density $\bar{p}(\theta) \propto \exp\left(-l(\theta) - \frac{\lambda_1}{\lambda_2}\|\theta\|^2\right)$ in \mathbb{R}^{d_θ} with $\|l\|_\infty \leq C$.

Therefore, we have $(\sigma(p_t^*))^2 \leq \frac{\lambda_2 m}{2\lambda_1} \exp\left(\frac{2}{\lambda_2}\right)$ and $(\sigma(q_t^*))^2 \leq \frac{\lambda_2 n}{2\lambda_1} \exp\left(\frac{6}{\lambda_2}\right)$. According to Theorem 22, we set the step size ι of ULA as

$$\iota \leq \frac{\lambda_1 \lambda_2 T^{-2J}}{32(3L_1 + \lambda_1)^2 \max\{m, n\}} \min \left\{ \frac{1}{2 \exp\left(\frac{12}{\lambda_2}\right)}, \frac{\lambda_1}{\lambda_2 \exp\left(\frac{24}{\lambda_2}\right)}, \frac{\lambda_1^2}{\lambda_2^2 \max\{m, n\} \exp\left(\frac{30}{\lambda_2}\right)} \right\},$$

and sub-problem iterations T_t at each step t

$$T_t \geq \iota^{-1} \frac{\lambda_2 e^{\frac{12}{\lambda_2}}}{\lambda_1} \left[\max \left\{ 3, \frac{24}{\lambda_2} + \log\left(\frac{\lambda_1}{\lambda_2}\right) + 2, \frac{30}{\lambda_2} + \log(\max\{m, n\}) + 3 \log\left(\frac{\lambda_1}{\lambda_2}\right) + 1 \right\} + 2J \log(T) + C \right],$$

to guarantee that $\max\{\delta_{t,1}, \delta_{t,2}\} \leq \frac{1}{2} \min\{c(p_t^*), c(q_t^*)\} T^{-2J}$ for any $t \geq 1$.

Under Corollary 9, we complete the proof. \blacksquare

Background on MALA and the proximal sampler: For simplicity, we use \mathcal{N} denote Gaussian distribution. Considering π be a target distribution on \mathbb{R}^{d_θ} , and denoting Q as a markov transition kernel, we have the Metropolis-Adjusted Langevin algorithm (Algorithm 2) with proposal kernel Q . In general, the target distribution $\pi \propto \exp\{-f\}$ and the proposal kernel is taken to be one step of the discretized Langevin algorithm, which is formulated as following:

$$Q(\theta, \cdot) = \mathcal{N}(\theta - \eta \nabla f(\theta), 2\eta \mathbf{I}_{d_\theta}). \quad (103)$$

Next, we introduce the proximal sampler (Algorithm 3) which has been mentioned in (Altschuler and Chewi, 2024). The third step of the Algorithm 3 approximates the objective distribution $\pi(\theta|y_k) \propto \exp\{-f(\theta) - \frac{1}{2\eta}\|\theta - y_k\|^2\}$ using MALA with high accuracy.

Algorithm 2 Metropolis-Adjusted Langevin Algorithm (MALA)

Input: number of outer iterations T , initial point θ_0 from a starting distribution π_0 , step size η
Output: sequence of samples $\theta_1, \dots, \theta_T$

- 1: **for** $t = 0$ to T **do**
 - 2: Propose $y_t \sim Q(\theta_t, \cdot)$.
 - 3: Compute the acceptance probability $\alpha_t = \min \left\{ 1, \frac{\pi(y_t)Q(y_t, \theta_t)}{\pi(\theta_t)Q(\theta_t, y_t)} \right\}$.
 - 4: Draw $u \sim \text{Unif}[0, 1]$.
 - 5: **if** $u < \alpha_t$ **then**
 - 6: Accept the proposal: $\theta_{t+1} \leftarrow y_t$.
 - 7: **else**
 - 8: Reject the proposal: $\theta_{t+1} \leftarrow \theta_t$.
 - 9: **end if**
 - 10: **end for**
-

Algorithm 3 Proximal Sampler with MALA

Input: number of outer iterations T , initial point θ_0 and step size η
Output: sequence of samples $\theta_1, \dots, \theta_T$

- 1: **for** $t = 0$ to T **do**
 - 2: Propose $y_t \sim \mathcal{N}(\theta_t, \eta \mathbf{I}_{d_\theta})$.
 - 3: Sample $\theta_{t+1} \sim \pi(\theta | y_t) \propto \exp\{-f(\theta) - \frac{1}{2\eta} \|\theta - y_t\|^2\}$ using Algorithm 2.
 - 4: **end for**
-

Proof [Proof of Corollary 14] According to the explicit expressions of p_t^* and q_t^* , we have that $h^{(t)}(\cdot)$ has $(3L_1 + \lambda_1)\lambda_2^{-1}$ -Lipschitz gradient and $g^{(t)}(\cdot)$ has $(L_1 + \lambda_1)\lambda_2^{-1}$ -Lipschitz gradient for any $t \geq 0$. Therefore, combining (Altschuler and Chewi, 2024, Theorem 5.3) and recalling that $p_t^* \propto \exp\{-g^{(t)}(x)\}$ satisfies log-Sobolev inequality with $\lambda_1 \exp\{-4\lambda_2^{-1}\}\lambda_2^{-1}$ and $q_t^* \propto \exp\{-h^{(t)}(y)\}$ satisfies log-Sobolev inequality with $\lambda_1 \exp\{-12\lambda_2^{-1}\}\lambda_2^{-1}$ for any $t \geq 0$ in the proof of Theorem 12, we set sub-problem iterations T_t at each step t

$$T_t \geq \tilde{\mathcal{O}} \left(\frac{(3L_1 + \lambda_1)e^{\frac{12}{\lambda_2}}}{\lambda_1} (\max\{m, n\})^{1/2} (2J \log(T) + 1 + \lambda_2^{-1})^4 \right), \quad (104)$$

to guarantee that $\max\{\delta_{t,1}, \delta_{t,2}\} \leq \frac{1}{2} \min\{c(p_t^*), c(q_t^*)\} T^{-2J}$ for any $t \geq 1$. ■

Appendix E. Stochastic Version

For simplicity, we denote $\psi_\xi(x) := \psi(x, \xi)$, $\phi_\xi(y) := \phi(y, \xi)$ for any random variable ξ and functions $\psi : \mathcal{X} \times \Lambda \rightarrow \mathbb{R}$, $\phi : \mathcal{Y} \times \Lambda \rightarrow \mathbb{R}$. Considering N i.i.d random variables (ξ_1, \dots, ξ_n) with $\xi_i = \xi$ for any $i \in [1 : N]$, we let $f_{\xi_N}(\theta) := \frac{1}{N} \sum_{i=1}^N f(\theta, \xi_i)$ for any $f : \Pi \times \Lambda \rightarrow \mathbb{R}$. In addition, we introduce some elementary symbols for $p(x) \in \mathcal{P}_2(\mathcal{X})$ satisfying LSI with

constant α_p and $q(y) \in \mathcal{P}_2(\mathcal{Y})$ satisfying LSI with constant α_q , and $\delta > 0$:

$$\begin{aligned} \text{ExpectErr}_p^x(\delta) &:= \left[\lambda_2 + \frac{2\lambda_1}{\alpha_p} \right] \delta + \left[32 + \frac{8\lambda_1\sigma(p)}{\sqrt{\alpha_p}} \right] \sqrt{\delta}, \\ \text{ExpectErr}_q^y(\delta) &:= \left[\lambda_2 + \frac{2\lambda_1}{\alpha_q} \right] \delta + \left[40 + \frac{8\lambda_1\sigma(q)}{\sqrt{\alpha_q}} \right] \sqrt{\delta}, \end{aligned}$$

We also let $\text{StoErr}_{\phi_N}^y(q) := \left| F_{\phi_{\bar{\xi}_N}}^y(q) - F_{\mathbb{E}_{\xi}\phi}^y(q) \right|$ and $\text{StoErr}_{\psi_N}^x(p) := \left| F_{\psi_{\bar{\xi}_N}}^x(p) - F_{\mathbb{E}_{\xi}\psi}^x(p) \right|$. To account for stochastic version, a modification of the functions $\hat{\phi}_t$ and $\hat{\psi}_t$ in the PAPAL algorithm iteration step for deterministic version is required. The PAPAL algorithm for stochastic version can be obtained by performing stochastic sampling on functions $\hat{\phi}_t$ and $\hat{\psi}_t$ (i.e. substituting $\hat{\phi}_t$ with $(\hat{\phi}_t)_{\bar{\xi}_{N,t}} := -(1 + \mu_t)\mathbb{E}_{\hat{p}_{t-1}} \left[G_{\bar{\xi}_{N,t}}(\cdot, y) \right] + \mu_t\mathbb{E}_{\hat{p}_{t-2}} \left[G_{\bar{\xi}_{N,t}}(\cdot, y) \right]$ and $\hat{\psi}_t$ with $(\hat{\psi}_t)_{\bar{\xi}_{N,t}} := \mathbb{E}_{\hat{q}_t} \left[G_{\bar{\xi}_{N,t}}(x, \cdot) \right]$) for each $t \in [1 : T]$ in Algorithm 1. In this context, $\bar{\xi}_{N,t}$ denotes the sampling of N i.i.d. random variables (ξ_1, \dots, ξ_N) associated with the t -th iteration. In the statement of the following lemma, we use function $(\hat{\phi}_t)_{\bar{\xi}_N}$ to denote $(\hat{\phi}_t)_{\bar{\xi}_{N,t}}$, and function $(\hat{\psi}_t)_{\bar{\xi}_N}$ to denote $(\hat{\psi}_t)_{\bar{\xi}_{N,t}}$ for consistency.

Lemma 32. *We set*

$$\gamma_t \mu_t = \gamma_{t-1}, \quad (105)$$

$$\gamma_t \tau_t \leq \gamma_{t-1}(\tau_{t-1} + \lambda_2), \quad (106)$$

$$\gamma_t \eta_t \leq \gamma_{t-1}(\eta_{t-1} + \lambda_2), \quad (107)$$

$$\tau_t \eta_{t-1} \geq \mu_t, \quad (108)$$

$$\mu_t \leq 1, \quad (109)$$

and assume that subproblem error bounds at each step- t satisfy

$$\text{KL}(p_t \| p_t^*) \leq \delta_{t,1}, \text{KL}(q_t \| q_t^*) \leq \delta_{t,2}, \quad (110)$$

then

$$\begin{aligned} \sum_{t=1}^k \gamma_t Q(w_t, w) &\leq \gamma_1 \eta_1 \text{KL}(p \| p_0^*) - \gamma_k (\eta_k + \lambda_2) \text{KL}(p \| p_k^*) + \gamma_1 \tau_1 \text{KL}(q \| q_0^*) \\ &\quad - \gamma_k \left(\tau_k + \lambda_2 - \frac{1}{4\eta_k} \right) \text{KL}(q \| q_k^*) \\ &\quad + \sum_{t=1}^k \gamma_t \left[4(1 + \mu_t)\epsilon + \text{ExpectErr}_{p_t^*}^x(\delta_{t,1}) + \text{ExpectErr}_{q_t^*}^y(\delta_{t,2}) \right] \\ &\quad + \sum_{t=1}^k \gamma_t \left[\text{StoErr}_{(\hat{\phi}_t)_n}^y(q_t^*) + \text{StoErr}_{(\hat{\phi}_t)_n}^y(q) + \text{StoErr}_{(\hat{\psi}_t)_n}^x(p_t^*) + \text{StoErr}_{(\hat{\psi}_t)_n}^x(p) \right], \end{aligned} \quad (111)$$

for any $\epsilon > 0$ and $p \in \mathcal{Q}_x(C_1), q \in \mathcal{Q}_y(C_2)$ with probability at least $1 - \delta$ for sample size $M = \Theta(\epsilon^{-2}[\log(T\delta^{-1}) + d \log(1 + L_0\sqrt{\lambda_2/\lambda_1}\epsilon^{-1})])$ in each iteration where $d = \max\{m, n\}$.

Moreover, we have

$$\begin{aligned}
 & \gamma_k \left(\tau_k + \lambda_2 - \frac{1}{4\eta_k} \right) \text{KL}(q_* \| q_k^*) \\
 & \leq \gamma_1 \eta_1 \text{KL}(p_* \| p_0^*) + \gamma_1 \tau_1 \text{KL}(q_* \| q_0^*) + \sum_{t=1}^k \gamma_t \left[4(1 + \mu_t) \epsilon + \text{ExpectErr}_{p_t^*}(\delta_{t,1}) + \text{ExpectErr}_{q_t^*}(\delta_{t,2}) \right] \\
 & \quad + \sum_{t=1}^k \gamma_t \left[\text{StoErr}_{(\hat{\phi}_t)_n}^y(q_t^*) + \text{StoErr}_{(\hat{\phi}_t)_n}^y(q_*) + \text{StoErr}_{(\hat{\psi}_t)_n}^x(p_t^*) + \text{StoErr}_{(\hat{\psi}_t)_n}^x(p_*) \right], \quad (112)
 \end{aligned}$$

and

$$\begin{aligned}
 & \gamma_k \left(\eta_k - \frac{1}{4(\tau_k + \lambda_2)} \right) \text{KL}(p_k^* \| p_{k-1}^*) \\
 & \leq \gamma_1 \eta_1 \text{KL}(p_* \| p_0^*) + \gamma_1 \tau_1 \text{KL}(q_* \| q_0^*) + \sum_{t=1}^k \gamma_t \left[4(1 + \mu_t) \epsilon + \text{ExpectErr}_{p_t^*}(\delta_{t,1}) + \text{ExpectErr}_{q_t^*}(\delta_{t,2}) \right] \\
 & \quad + \sum_{t=1}^k \gamma_t \left[\text{StoErr}_{(\hat{\phi}_t)_n}^y(q_t^*) + \text{StoErr}_{(\hat{\phi}_t)_n}^y(q_*) + \text{StoErr}_{(\hat{\psi}_t)_n}^x(p_t^*) + \text{StoErr}_{(\hat{\psi}_t)_n}^x(p_*) \right] \quad (113)
 \end{aligned}$$

with respect to the MNE $w_* = (p_*, q_*)$ of Eq. (23).

Proof According to optimality condition (Lemma 30), we have

$$F_{(\hat{\phi}_t)_{\bar{\xi}_n}}^y(q_t^*) - F_{(\hat{\phi}_t)_{\bar{\xi}_n}}^y(q) = \tau_t [\text{KL}(q \| q_{t-1}^*) - \text{KL}(q_t^* \| q_{t-1}^*)] - (\tau_t + \lambda_2) \text{KL}(q \| q_t^*), \quad (114)$$

$$F_{(\hat{\psi}_t)_{\bar{\xi}_n}}^y(p_t^*) - F_{(\hat{\psi}_t)_{\bar{\xi}_n}}^y(p) = \eta_t [\text{KL}(p \| p_{t-1}^*) - \text{KL}(p_t^* \| p_{t-1}^*)] - (\eta_t + \lambda_2) \text{KL}(p \| p_t^*). \quad (115)$$

By using Eq. (114), we have

$$\begin{aligned}
 F_{\mathbb{E}_\xi \hat{\phi}_t}^y(q_t) - F_{\mathbb{E}_\xi \hat{\phi}_t}^y(q) &= F_{\mathbb{E}_\xi \hat{\phi}_t}^y(q_t) - F_{(\hat{\phi}_t)_{\bar{\xi}_n}}^y(q_t) + F_{(\hat{\phi}_t)_{\bar{\xi}_n}}^y(q) - F_{\mathbb{E}_\xi \hat{\phi}_t}^y(q) \\
 & \quad + \underbrace{F_{(\hat{\phi}_t)_{\bar{\xi}_n}}^y(q_t^*) - F_{(\hat{\phi}_t)_{\bar{\xi}_n}}^y(q) + F_{(\hat{\phi}_t)_{\bar{\xi}_n}}^y(q_t) - F_{(\hat{\phi}_t)_{\bar{\xi}_n}}^y(q_t^*)}_a \\
 & \stackrel{\text{b}}{\leq} \text{StoErr}_{(\hat{\phi}_t)_n}^y(q_t^*) + \text{StoErr}_{(\hat{\phi}_t)_n}^y(q) + \tau_t [\text{KL}(q \| q_{t-1}^*) - \text{KL}(q_t^* \| q_{t-1}^*)] \\
 & \quad - (\tau_t + \lambda_2) \text{KL}(q \| q_t^*) + \left[\lambda_2 + \frac{2\lambda_1}{\alpha_{q_t^*}} \right] \delta_{t,2} + \left[12 + \frac{8\lambda_1 \sigma(q_t^*)}{\sqrt{\alpha_{q_t^*}}} \right] \sqrt{\delta_{t,2}}, \quad (116)
 \end{aligned}$$

where (b) is derived from applying Eq.(80) to (a) and the estimation that

$$\left| \text{StoErr}_{(\hat{\phi}_t)_n}^y(q_t^*) - \text{StoErr}_{(\hat{\phi}_t)_n}^y(q_t) \right| \leq 2\sqrt{\delta_{t,2}},$$

which follows from Pinsker's inequality directly. Likewise, we can deduce that

$$\begin{aligned}
 F_{\mathbb{E}_\xi \hat{\psi}_t}^x(p_t) - F_{\mathbb{E}_\xi \hat{\psi}_t}^x(p) &\leq \text{StoErr}_{(\hat{\psi}_t)_n}^x(p_t^*) + \text{StoErr}_{(\hat{\psi}_t)_n}^x(p) + \eta_t [\text{KL}(p||p_{t-1}^*) - \text{KL}(p_t^*||p_{t-1}^*)] \\
 &\quad - (\eta_t + \lambda_2)\text{KL}(p||p_t^*) + \left[\lambda_2 + \frac{2\lambda_1}{\alpha_{p_t^*}} \right] \delta_{t,1} + \left[4 + \frac{8\lambda_1\sigma(p_t^*)}{\sqrt{\alpha_{p_t^*}}} \right] \sqrt{\delta_{t,1}}.
 \end{aligned} \tag{117}$$

By utilizing the boundedness of f and carefully selecting constants C_1, C_2 , we can adjust the number of sampling points for both \hat{p}_t and \hat{q}_t to obtain the following result

$$|\text{BiLin}_{\mathbb{E}_\xi f}(p_t - \hat{p}_t, q)| \leq \epsilon, \quad |\text{BiLin}_{\mathbb{E}_\xi f}(p, q_t - \hat{q}_t)| \leq \epsilon, \tag{118}$$

with probability $1 - \delta$ for all $q(y) \in \mathcal{Q}_Y(C_2)$, $p(x) \in \mathcal{Q}_X(C_1)$ uniformly (Refer to Lemma 7 in Appendix F). In order to deal with finite particles, we denote $(\mathbb{E}_\xi \phi_t)(y) = -(1 + \mu_t)\mathbb{E}_{p_{t-1}}[\mathbb{E}_\xi f(x, y, \xi)] + \mu_t\mathbb{E}_{p_{t-2}}[\mathbb{E}_\xi f(x, y, \xi)]$, $(\mathbb{E}_\xi \psi_t)(x) = \mathbb{E}_{q_t}[\mathbb{E}_\xi f(x, y, \xi)]$. Similar to the proof of Theorem 31, we leverage the definition of the Max-Min gap function Q to arrive at

$$\begin{aligned}
 \sum_{t=1}^k \gamma_t Q(w_t, w) &\leq \sum_{t=1}^k \gamma_t [\eta_t \text{KL}(p||p_{t-1}^*) - (\eta_t + \lambda_2)\text{KL}(p||p_t^*)] \\
 &\quad + \sum_{t=1}^k \gamma_t [\tau_t \text{KL}(q||q_{t-1}^*) - (\tau_t + \lambda_2)\text{KL}(q||q_t^*)] + 4 \sum_{t=1}^k \gamma_t (1 + \mu_t) \epsilon \\
 &\quad + \sum_{t=1}^k \gamma_t [\mu_t \text{BiLin}_{\mathbb{E}_\xi f}(p_{t-1}^* - p_{t-2}^*, q_{t-1}^* - q) - \text{BiLin}_{\mathbb{E}_\xi f}(p_t^* - p_{t-1}^*, q_t^* - q)] \\
 &\quad - \sum_{t=1}^k \gamma_t [\tau_t \text{KL}(q_t^*||q_{t-1}^*) + \eta_t \text{KL}(p_t^*||p_{t-1}^*) + \mu_t \text{BiLin}_{\mathbb{E}_\xi f}(p_{t-1}^* - p_{t-2}^*, q_{t-1}^* - q_t^*)] \\
 &\quad + \sum_{t=1}^k \gamma_t \left[\text{StoErr}_{(\hat{\phi}_t)_n}^y(q_t^*) + \text{StoErr}_{(\hat{\phi}_t)_n}^y(q) + \text{StoErr}_{(\hat{\psi}_t)_n}^x(p_t^*) + \text{StoErr}_{(\hat{\psi}_t)_n}^x(p) \right] \\
 &\quad + \sum_{t=1}^k \gamma_t \left[\text{ExpectErr}_{p_t^*}^x(\delta_{t,1}) + \text{ExpectErr}_{q_t^*}^x(\delta_{t,2}) \right].
 \end{aligned} \tag{119}$$

Therefore, it's direct to obtain

$$\begin{aligned}
 \sum_{t=1}^k \gamma_t Q(w_t, w) &\leq \gamma_1 \eta_1 \text{KL}(p||p_0^*) - \gamma_k (\eta_k + \lambda_2) \text{KL}(p||p_k^*) + \gamma_1 \tau_1 \text{KL}(q||q_0^*) \\
 &\quad - \gamma_k (\tau_k + \lambda_2) \text{KL}(q||q_k^*) - \gamma_k \eta_k \text{KL}(p_k^*||p_{k-1}^*) - \gamma_k \text{BiLin}_{\mathbb{E}_\xi f}(p_k^* - p_{k-1}^*, q_k^* - q) \\
 &\quad + \sum_{t=1}^k \gamma_t \left[\text{StoErr}_{(\hat{\phi}_t)_n}^y(q_t^*) + \text{StoErr}_{(\hat{\phi}_t)_n}^y(q) + \text{StoErr}_{(\hat{\psi}_t)_n}^x(p_t^*) + \text{StoErr}_{(\hat{\psi}_t)_n}^x(p) \right] \\
 &\quad + \sum_{t=1}^k \gamma_t \left[4(1 + \mu_t) \epsilon + \text{ExpectErr}_{p_t^*}^x(\delta_{t,1}) + \text{ExpectErr}_{q_t^*}^x(\delta_{t,2}) \right].
 \end{aligned} \tag{120}$$

In light of the prior analysis, we can complete the proof of our lemma by following the same steps outlined in Theorem 31. \blacksquare

By utilizing Assumptions **[A₄]**, **[A₅]** and **[A₆]**, and applying $\hat{\phi}_t(y, \xi)$ and $\hat{\psi}_t(x, \xi)$ to Lemma 7, we can derive that the discretization error of finite particles on Λ satisfies

$$\text{StoErr}_{(\hat{\phi}_t)_N}^y(q) \leq \epsilon, \text{StoErr}_{(\hat{\psi}_t)_N}^x(p) \leq \epsilon, \quad (121)$$

with probability $1 - \delta$ for any $p \in \mathcal{Q}_x(C_1)$ and $q \in \mathcal{Q}_y(C_2)$, where the required number of particles is

$$N = \mathcal{O} \left(\epsilon^{-2} \left[-\log(\delta) + d \log \left(1 + L_0 r_{\max\{C_1, C_2\}}(\epsilon, \lambda_1, \lambda_2) \epsilon^{-1} \right) \right] \right), \quad (122)$$

and $r_C(\epsilon, \lambda_1, \lambda_2) := 8\sqrt{\frac{\lambda_2 + C}{\lambda_1} \max\{\log(\epsilon^{-1}), m, n\}}$. Therefore, combining Lemma 32 and aforementioned discretization error of finite particles, we obtain stochastic outer loop convergence result Lemma 15. Applying inner loop error bound to Lemma 15, we can derive the global convergence result Theorem 16 and Corollary 17 w.r.t stochastic PAPAL.

Appendix F. Discretization Error of Finite Particles

Proof [Proof of Lemma 7] In order to improve the overall clarity of our argument, we concentrate on establishing the estimation in Eq. (17). We split the proof into two parts.

Part I Combing Lemma 23, and the definition of $\mathcal{Q}_y(C_2)$, we have

$$\begin{aligned} \left| \iint_{\mathbb{R}^m \times \{y \mid \|y\| \geq r_y\}} (\hat{p}_t - p_t)(x) f(x, y) q(y) dy dx \right| &\leq 2 \exp\left(\frac{2C_2}{\lambda_2}\right) \mathbb{P}[\|\hat{y}\|^2 \geq r_y^2] \\ &\leq 2 \exp\left(\frac{2C_2}{\lambda_2}\right) \exp\left(-\frac{\lambda_1 r_y^2}{10\lambda_2}\right), \end{aligned} \quad (123)$$

for any $q \in \mathcal{Q}_y(C_2)$ with $r_y^2 \geq \frac{n\lambda_2}{\lambda_1}$ where $\hat{y} \sim \mathcal{N}\left(0, \frac{\lambda_2}{2\lambda_1} I_n\right)$. Then, we derive the bound:

$$\left| \iint_{\mathbb{R}^m \times \{y \mid \|y\| \geq r_y\}} (\hat{p}_t - p_t)(x) f(x, y) q(y) dy dx \right| \leq \frac{1}{4} \epsilon$$

for any $q \in \mathcal{Q}_y(C_2)$ by setting $r_y \geq 4\sqrt{\frac{\lambda_2}{\lambda_1} \max\{\log(\epsilon^{-1}), n\} + 2\frac{\lambda_2}{\lambda_1} + 2\frac{C_2}{\lambda_1}}$.

Part II. We need to consider the uniform bound between \hat{p}_t and p_t when $y \in \mathbb{B}_{r_y}(\mathbf{0})$ with $r_y = 8\sqrt{\frac{\lambda_2 + C_2}{\lambda_1} \max\{\log(\epsilon^{-1}), n\}}$. According to Lemma 5.7 in Wainwright (2019), we have the Δ -covering number (i.e. $N(\Delta; \mathbb{B}_{r_y}(\mathbf{0}), \|\cdot\|)$) of $\mathbb{B}_{r_y}(\mathbf{0}) := \{y \in \mathbb{R}^n \mid \|y\| \leq r_y\}$ in the $\|\cdot\|$ -norm obeys the bound

$$N(\Delta; \mathbb{B}_{r_y}(\mathbf{0}), \|\cdot\|) \leq \left(1 + \frac{2r_y}{\Delta}\right)^n. \quad (124)$$

On the other hand, we obtain

$$\mathbb{P} \left[\left| \int_{\mathbb{R}^m} (\hat{p}_t - p_t)(x) f(x, y) dx \right| > \rho \right] \leq 2 \exp \left(-\frac{\rho^2 M}{2} \right), \quad (125)$$

under Lemma 24 when fixing $y \in \mathbb{B}_{r_y}(\mathbf{0})$. Therefore, we can deduce the concentration inequality on a Δ -cover $\mathcal{C}_{\mathbb{B}_{r_y}(\mathbf{0}), \Delta}$ of $\mathbb{B}_{r_y}(\mathbf{0})$

$$\mathbb{P} \left[\max_{y \in \mathcal{C}_{\mathbb{B}_{r_y}(\mathbf{0}), \Delta}} \left| \int_{\mathbb{R}^m} (\hat{p}_t - p_t)(x) f(x, y) dx \right| > \rho \right] \leq 2 \left(1 + \frac{2r_y}{\Delta} \right)^n \exp \left(-\frac{\rho^2 M}{2} \right), \quad (126)$$

by combining Eq. (124) and Eq. (125). Then, the following estimation

$$\mathbb{P} \left[\max_{y \in \mathbb{B}_{r_y}(\mathbf{0})} \left| \int_{\mathbb{R}^m} (\hat{p}_t - p_t)(x) f(x, y) dx \right| > \rho + L_0 \Delta \right] \leq 2 \left(1 + \frac{2r_y}{\Delta} \right)^n \exp \left(-\frac{\rho^2 M}{2} \right), \quad (127)$$

follows from Assumption **[A₂]**, Eq. (126) and Lemma.

Finally, we conclude the result in (17) by using **Part I** and **Part II** with setting $\Delta = \frac{\epsilon}{8L_0}$, $\rho = \frac{\epsilon}{8}$ and $M \geq 128 \left[-\log(\delta/2) + n \log \left(1 + \frac{16r_y L_0}{\epsilon} \right) \right] \epsilon^{-2}$ and the result in (17) can be derived under the similar argument as well. ■

Appendix G. Empirical Results

In this section, we would like to illustrate the motivation of our algorithm compared to weight-driven algorithm – WFR-DA (Domingo-Enrich et al., 2020). Thus, the theoretical guarantees for particle-based algorithm is important in real practice.

Non-Convex Optimization. Our algorithm can be recognized as a generalization of non-convex optimization. When the $f(x, y)$ is a constant with respect to y , it is a minimization problem. Our first example is a non-convex minimization task, which provides us insights to distinguish the optimization on weights and the positions. Figure 1 demonstrates that the weight-driven algorithms only “select” particles close to the solution, while particle-based algorithms move the particles to find the solution. The difference makes weight-driven algorithms need more particles to make sure that the initialized some particles can approximate the solution. Unless the sample size is extremely large, the weight-driven algorithm cannot obtain good solution. In fact, considering the volume of space, the sample size of weight-driven algorithms is exponential with respect to the dimension, which makes the original algorithm (with theoretical guarantees) in (Domingo-Enrich et al., 2020) hardly usable in high dimensions. It is worth mentioning that the combination of weight optimization and the particle-based optimization might be better heuristically, but the leading driven force of the dynamics should be the particle-based one. The theoretical analysis for the combined algorithm mainly driven by particle-based force would be a good future work.

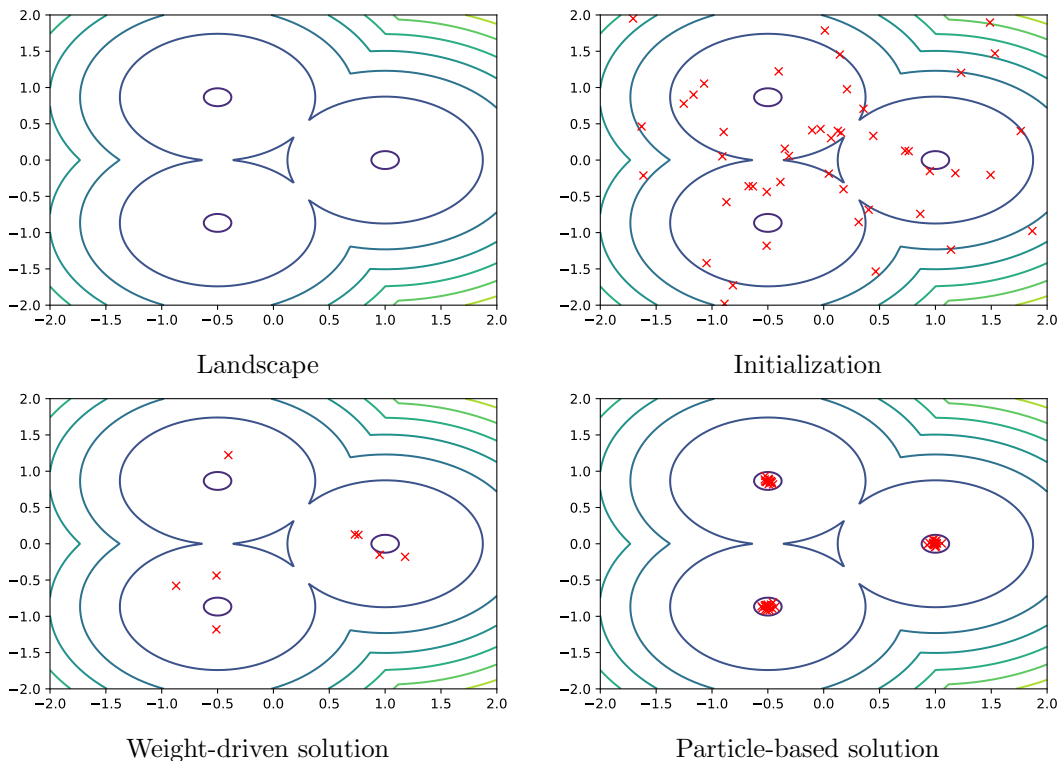


Figure 1: Comparison between weight-driven and particle-based algorithms.

Generative Adversarial Networks. We consider the training process of Generative Adversarial Networks to illustrate the comparison between weight-driven and particle-based algorithms. In particular, when the dimension is large, the weight-driven algorithm would suffer from the curse of dimensionality. For simplicity, we choose the Gaussian target distribution and compute the KL divergence between generated and target distribution. We choose a randomly sampled mean as our target distribution and use a linear model to reproduce it. We choose the step-size 10^{-3} to perform the in our experiments. Figure 2 indicate that the weight-driven algorithms (WFR-DA) suffer from the curse of dimensionality seriously, while the particle-based algorithm can perform well on high dimension space. The empirical results have justify the importance of the particle-based algorithm and efficacy of PAPAL in high dimension spaces.




Figure 2: Illustration of the sample efficiency on Generative Adversarial Networks. (x -axis is the dimension; y -axis is the KL divergence to measure the distance between real and fake distribution; n denotes the number of particles)